DOCUMENT RESUME

ED 379 317 TM 022 682

AUTHOR	Forbus, Kenneth D.; And Others
TITLE	MAC/FAC: A Model of Similarity-Based Retrieval. Technical Report #59.
INSTITUTION	Northwestern Univ., Evanston, IL. Inst. for the Learning Sciences.
SPONS AGENCY	Office of Naval Research, Arlington, Va.
PUB DATE	Oct 94
CONTRACT	N00014-89-J-1272
NOTE	68p.; Type in distribution list is very small.
PUB TYPE	Reports - Evaluative/Feasibility (142)
EDRS PRICE	MF01/PC03 Plus Postage.
DESCRIPTORS	*Data Analysis; Decision Making; *Encoding (Psychology); Models; Psychological Studies; *Recall (Psychology); Retention (Psychology)
IDENTIFIERS	*MAC FAC Model; *Retrieval (Memory); Similarity (Concept)

ABSTRACT

A model of similarity-based retrieval is presented that attempts to capture these seemingly contradictory psychological phenomena: (1) structural commonalities are weighed more heavily than surface commonalities in soundness or similarity judgments (when both members are present); (2) superficial similarity is more important in retrieval from long-term memory than is structural similarity; and yet (3) purely structural (analogical) remindings are sometimes experienced. The model, called MAC/FAC (for "many are called, but few are chosen") consists of a first stage (MAC) that uses a computationally cheap, nonstructured matcher to filter candidates from a pool of memory items. Structured representations are redundantly encoded as "content vectors," whose dot product yields an estimate of how well the corresponding structural representations will match. The second stage (FAC) uses the Structure-Mapping Engine to compute a true structural match between the probe and output from the first stage. MAC/FAC has been fully implemented and tested on dozens of examples, and is shown to be capable of modeling patterns of access in psychological data. The relationship to other models of memory is discussed. Seven figures and 19 tables. (Contains 128 references.) (Author/SLD)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

The Institute for the Learning Sciences

Northwestern University

TM

ED 379 317

MAC/FAC: A MODEL OF SIMILARITY-BASED RETRIEVAL

Kenneth D. Forbus
Dedre Gentner
Keith Law

Technical Report #59 • October 1994

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

- This document has been reproduced as received from the person or organization originating it.
- ☐ Minor changes have been made to improve reproduction quality.

- Points of view or opinions stated in this document do not necessarily represent official OERI position or policy.

TM022682

Established in 1989 with the support of Andersen Consulting

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
		Technical Report 9/1/89 - 10/31/92

4. TITLE AND SUBTITLE
MAC/FAC: A model of similarity-based retrieval

5. FUNDING NUMBERS
C - N00014-89-J-1272

6. AUTHOR(S)
Kenneth D. Forbus & Dedre Gentner

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
The Institute for the Learning Sciences
Northwestern University
1890 Maple Avenue
Evanston, IL 60201

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Cognitive Science Program (1142CS)
Office of Naval Research
800 N. Quincy Street
Arlington, VA 22217-5000

10. SPONSORING / MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION / AVAILABILITY STATEMENT
Approved for Public Release
Unlimited distribution

12b. DISTRIBUTION CODE

13. ABSTRACT (Maximum 200 words) We present a model of similarity-based retrieval that attempts to capture these seemingly contradictory psychological phenomena: (1) structural commonalities are weighed more heavily than surface commonalities in soundness or similarity judgments (when both members are present); (2) superficial similarity is more important in retrieval from long-term memory than is structural similarity; and yet (3) purely structural (analogical) remindings are sometimes experienced. Our model, called MAC/FAC (for "many are called but few are chosen") consists of two stages. The first stage (MAC) uses a computationally cheap, non-structured matcher to filter candidates from a pool of memory items. We redundantly encode structured representations as content vectors, whose dot product yields an estimate of how well the corresponding structural representations will match. The second stage (FAC) uses SME to compute a true structural match between the probe and output from the first stage. MAC/FAC has been fully implemented and tested on dozens of examples. We show that MAC/FAC is capable of modeling patterns of access found in psychological data, and illustrative via sensitivity analyses that these results exhibit the desired dependence on theoretically important factors. The relationship of MAC/FAC to other models of memory is discussed, along with implications and possible extensions.

14. SUBJECT TERMS	15. NUMBER OF PAGES 69
	16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UL	UL	UL	UL

NSN 7540-01-280-5500

Standard Form 298 (Rev 2-89)
Prescribed by ANSI Std Z39-18
298-102

GENERAL INSTRUCTIONS FOR COMPLETING SF 298

The Report Documentation Page (RDP) is used in announcing and cataloging reports. It is important that this information be consistent with the rest of the report, particularly the cover and title page. Instructions for filling in each block of the form follow. It is important to ***stay within the lines*** to meet ***optical scanning requirements***.

Block 1. Agency Use Only *(Leave blank)*.

Block 2. Report Date. Full publication date including day, month, and year, if available (e.g. 1 Jan 88). Must cite at least the year.

Block 3. Type of Report and Dates Covered. State whether report is interim, final, etc. If applicable, enter inclusive report dates (e.g. 10 Jun 87 - 30 Jun 88).

Block 4. Title and Subtitle. A title is taken from the part of the report that provides the most meaningful and complete information. When a report is prepared in more than one volume, repeat the primary title, add volume number, and include subtitle for the specific volume. On classified documents enter the title classification in parentheses.

Block 5. Funding Numbers. To include contract and grant numbers; may include program element number(s), project number(s), task number(s), and work unit number(s). Use the following labels:

C	- Contract	**PR**	- Project
G	- Grant	**TA**	- Task
PE	- Program Element	**WU**	- Work Unit Accession No.

Block 6. Author(s). Name(s) of person(s) responsible for writing the report, performing the research, or credited with the content of the report. If editor or compiler, this should follow the name(s).

Block 7. Performing Organization Name(s) and Address(es). Self-explanatory.

Block 8. Performing Organization Report Number. Enter the unique alphanumeric report number(s) assigned by the organization performing the report.

Block 9. Sponsoring/Monitoring Agency Name(s) and Address(es). Self-explanatory.

Block 10. Sponsoring/Monitoring Agency Report Number. *(If known)*

Block 11. Supplementary Notes. Enter information not included elsewhere such as: Prepared in cooperation with...; Trans. of...; To be published in.... When a report is revised, include a statement whether the new report supersedes or supplements the older report.

Block 12a. Distribution/Availability Statement. Denotes public availability or limitations. Cite any availability to the public. Enter additional limitations or special markings in all capitals (e.g. NOFORN, REL, ITAR).

- **DOD** - See DoDD 5230.24, "Distribution Statements on Technical Documents."
- **DOE** - See authorities.
- **NASA** - See Handbook NHB 2200.2.
- **NTIS** - Leave blank.

Block 12b. Distribution Code.

- **DOD** - Leave blank.
- **DOE** - Enter DOE distribution categories from the Standard Distribution for Unclassified Scientific and Technical Reports.
- **NASA** - Leave blank.
- **NTIS** - Leave blank.

Block 13. Abstract. Include a brief *(Maximum 200 words)* factual summary of the most significant information contained in the report.

Block 14. Subject Terms. Keywords or phrases identifying major subjects in the report.

Block 15. Number of Pages. Enter the total number of pages.

Block 16. Price Code. Enter appropriate price code *(NTIS only)*.

Blocks 17. - 19. Security Classifications. Self-explanatory. Enter U.S. Security Classification in accordance with U.S. Security Regulations (i.e., UNCLASSIFIED). If form contains classified information, stamp classification on the top and bottom of the page.

Block 20. Limitation of Abstract. This block must be completed to assign a limitation to the abstract. Enter either UL (unlimited) or SAR (same as report). An entry in this block is necessary if the abstract is to be limited. If blank, the abstract is assumed to be unlimited.

MAC/FAC: A Model of Similarity-based Retrieval

Kenneth D. Forbus
The Institute for the Learning Sciences

Dedre Gentner
Psychology Department

Keith Law
The Institute for the Learning Sciences

Northwestern University
1890 Maple Avenue,
Evanston, IL, 60201

Reproduction in whole or part is permitted for any purpose of the United States Government.

This research was sponsored by the Cognitive Science Program, Office of Naval Research, under Grant No. N00014-89-J-1272, Contract Authority Identification No. NR-442f-007.

Approved for public release; distribution unlimited.

ABSTRACT: We present a model of similarity-based retrieval that attempts to capture three seemingly contradictory psychological phenomena: (1) structural commonalities are weighed more heavily than surface commonalities in similarity judgments for items in working memory; (2) in retrieval, superficial similarity is more important than structural similarity; and yet (3) purely structural (analogical) remindings are sometimes experienced. Our model, MAC/FAC, explains these phenomena in terms of a two-stage process. The first stage uses a computationally cheap, non-structural matcher to filter candidate long-term memory items. It uses *content vectors*, a redundant encoding of structured representations whose dot product estimates how well the corresponding structural representations will match. The second stage uses SME to compute structural matches on the handful of items found by the first stage. We show the utility of the MAC/FAC model through a series of computational experiments: (1) We demonstrate that MAC/FAC can model patterns of access found in psychological data, (2) We argue via sensitivity analyses that these simulation results rely on the theory, and (3) We compare the performance of MAC/FAC with ARCS, an alternate model of similarity-based retrieval, and demonstrate that MAC/FAC explains the data better than ARCS. Finally, we discuss limitations and possible extensions of the model, relationships with other recent retrieval models, and place MAC/FAC in the context of other recent work on the nature of similarity.

1. Introduction

Similarity-based remindings range from the sublime to the stupid. At one extreme, seeing the periodic table of elements reminds one of octaves in music. At the other, a bicycle reminds one of a pair of eyeglasses. Often, remindings are neither brilliant nor superficial but simply mundane, as when a bicycle reminds one of another bicycle. Theoretical attention is inevitably drawn to spontaneous analogy: i.e., to structural similarity unsupported by surface similarity, as in the octave/periodic table comparison. Such remindings seem clearly insightful and seem linked to the creative process, and should be included in any model of retrieval. But as we review below, research on the psychology of memory retrieval points to a preponderance of the latter two types of similarity -- (mundane) literal similarity, based on both structural and superficial commonalities -- and (dumb) superficial similarity, based on surface commonalities. A major challenge for research on similarity-based reminding is to devise a model that will produce chiefly literal-similarity and superficial remindings, but still produce occasional analogical remindings.

A further constraint on models of access comes from considering the role of similarity in transfer and inference. The large number of superficial remindings indicates that retrieval is not very sensitive to structural soundness. But appropriate transfer requires structural soundness, so that knowledge can be exported from one description into another. And psychological evidence (also discussed below) indicates that the mapping process involved in transfer is actually very sensitive to structural soundness. Hence our memories often give us information we don't want, which at first seems somewhat paradoxical. Any model of retrieval should explain this paradox.

This paper presents MAC/FAC, a model of similarity-based reminding that attempts to capture these phenomena. MAC/FAC models similarity-based retrieval as a two-stage process. The first stage (MAC) uses a cheap, non-structural matcher to quickly filter potentially relevant items from a pool of such items. These potential matches are then processed in the FAC stage by a more powerful (but more expensive) structural matcher, based on the structure-mapping notion of literal similarity (Gentner, 1983).

We begin in Section 2 by briefly reviewing psychological evidence on similarity-based retrieval and mapping, thereby extracting some criteria which retrieval models must satisfy. This section also outlines the computational issues raised by similarity-based retrieval, drawing on the AI literature as necessary. Section 3 describes the MAC/FAC model, showing how it satisfies the psychological and computational desiderata. 4 illustrates the model's psychological plausibility by simulating the results of a psychological experiment. Section 5 explores the consequences of different design decisions by sensitivity analyses at the level of algorithms, demonstrating that the model's performance depends on the theoretically important parameters. Section 6 compares MAC/FAC with ARCS, the closest competing model of similarity-based retrieval, demonstrating that MAC/FAC performs well on databases designed by others (e.g., the ARCS datasets) and that MAC/FAC's performance fits the psychological evidence better than ARCS. Finally, Section 7 compares MAC/FAC to several other memory models, analyzes some of its limitations, and discusses possible extensions.

2. Framework

Similarity-based transfer can be decomposed into subprocesses. Given that a person has some current *target* situation in working memory, transfer from prior knowledge requires at least

1. *Accessing* a similar (*base*) situation in long-term memory,
2. *Creating a mapping* from the base to the target, and
3. *Evaluating* the mapping.

In this case the base is an item from memory, and the target is the probe; that is, we think of the retrieved memory items as mapped to the probe. Other processes may also occur -- *verifying new inferences* about the target (Clement, 1986), *elaborating* the base and target (Ross, 1987; Falkenhainer, 1988), *adapting* or *tweaking* the domain representations to improve the match (Falkenhainer, 1990; Holyoak & Novick, in press; Kass, 1986, 1989), and *abstracting* the common structure from base and target (Gick & Holyoak, 1983; Skorstad, Gentner & Medin, 1988; Winston, 1982) -- but our focus is on the first three processes.

2.1 Structure-Mapping and the typology of similarity

The process of *mapping* aligns two representations and uses this alignment to generate analogical inferences (Gentner, 1983, 1988, 1989a). Alignment occurs via matching, which creates correspondences between items in the two representations. Analogical inferences are generated by using the correspondences to import knowledge from the base representation into the target. The mapping process is assumed to be governed by the constraints of *structural consistency*: *one-to-one mapping* and *parallel connectivity*. *One-to-one mapping* means that an interpretation of a comparison cannot align (e.g., place into correspondence) the same item in the base with multiple items in the target, or vice-versa. *Parallel connectivity* means that if an interpretation of a comparison aligns two statements, their arguments must also be placed into correspondence.[1] In this account, similarity is defined in terms of correspondences between structured representations (Gentner 1983; Gentner & Markman, 1993, 1994, in press; Goldstone, Medin & Gentner, 1991; Markman & Gentner, 1990,1993a,b; Medin, Goldstone & Gentner, 1993). Matches can be distinguished according to the kinds of commonalities present. An *analogy* is a match based on a common system of relations, especially involving higher-order relations.[2] A *literal similarity* match includes both common relational structure and common object descriptions. A *surface similarity* or *mere-appearance* match is based primarily on common object descriptions, with perhaps a few shared first-order relations.

There is considerable evidence that the mapping process is sensitive to structural commonalities. People can readily align two situations, preserving structurally important commonalties, making the appropriate lower-order substitutions, and mapping additional predicates into the target as *candidate inferences*. For example, Clement & Gentner (1991) showed people analogies and asked which of two lower-order assertions, both shared by base and target, was most important to the match. Subjects chose assertions that were connected to matching causal antecedents: that is, their choice was based not only on the goodness of the local match but on whether it was connected to a larger matching system. In a second study, subjects were asked to make a new prediction about the target based on the analogy with the base story. They again showed sensitivity to connectivity and systematicity in choosing which predicates to map as candidate inferences from base to target. Evidence for structural consistency in mapping comes from a study by Spellman and Holyoak (1992). They asked people to explicate the analogy between the Gulf War and World War II, assuming Saddam Hussein maps onto Hitler. Although people were divided in their mappings, they were highly consistent. People who mapped Bush onto Churchill mapped the current USA onto WWII Britain and people who mapped Bush onto FDR mapped the USA today onto USA during World War II.

---

[1] Previously we used the term *structurally grounded* for parallel connectivity.

[2] We define the *order* of an item in a representation as follows: Objects and constants are order 0. The order of a statement is one plus the maximum of the order of its arguments.

The degree of relational match is also important in determining people's evaluations of comparisons. People rate metaphors as more apt when they are based on relational commonalities than when they are based on common object-descriptions (Gentner, 1988b; Gentner & Clement, 1988). Gentner, Rattermann and Forbus (1993) asked subjects to rate the soundness and similarity of story pairs that varied in which kinds of commonalities they shared. Subjects' soundness and similarity ratings were substantially greater for pairs that shared higher-order relational structure than for those that did not (Gentner & Landers, 1985; Gentner, Rattermann, & Forbus, in press; Rattermann & Gentner, 1987). Common relational structure also contributes strongly to judgments of perceptual similarity (Goldstone, Medin & Gentner, 1991) as well as to the way in which people align pairs of pictures in a mapping task (Markman & Gentner, 1990, 1993b) and determine common and distinctive features (Gentner & Markman, 1994; Markman & Gentner, 1993a).

Any model of human similarity and analogy must capture this sensitivity to structural commonality. To do so, it must involve structural representations and processes that operate to align them (Barnden, 1994; Gentner & Markman, in press; Goldstone, Medin & Gentner, 1991; Holyoak, Novick, & Melz, 1994; Keane, 1988ab; Markman & Gentner, 1993; Medin, Goldstone, & Gentner, 1993; Reed, 1987; Reeves & Weisberg, 1994). This would seem to require abandoning some highly influential models of similarity: e.g., modeling similarity as the intersection of independent feature sets or as the dot product of feature vectors. However, we show below that a variant of these nonstructural models can be useful in describing memory retrieval.

2.1.1 Similarity-based Access from Long-term Memory

There is considerable evidence that access to long-term memory relies more on surface commonalities and less on structural commonalities than does mapping. For example, people often fail to access potentially useful analogs, as in Gick and Holyoak's (1980, 1983) dramatic demonstration. When subjects were told a story and then given an analogous problem to solve, about 30% solved the problem. However, if subjects were simply told to think about the story they had heard, 80% --- 90% solved the problem. We can infer that most of the subjects retained representations of the prior story sufficient to provide a useful analogy, but that hearing the structurally analogous problem did not provide spontaneous access to the story representation in memory. Other research has shown that, although people in a problem-solving task are often reminded of prior problems, these remindings are often based on surface similarity rather than on structural similarities between the solution principles (Holyoak & Koh, 1987; Keane, 1987, 1988b; Novick 1988; Reed, Ernst, & Banerji, 1974; Ross, 1984, 1987, 1989; see also the comprehensive review by Reeves & Weisberg, 1994).

The experiments we will model here investigated which kinds of similarities led to the best retrieval from long-term memory (Gentner & Landers, 1985; Gentner, Rattermann, & Forbus, 1993; Rattermann & Gentner, 1987). Subjects were first given a relatively large memory set (the "Karla the hawk" stories). About a week later, subjects were given new stories that resembled the original stories in various ways and were asked to write out any remindings they experienced to the prior stories while reading the new stories. Finally, they rated all the pairs for soundness -- i.e., how well inferences could be carried from one story to the other. The results showed a marked disassociation between retrieval and subjective soundness and similarity. Surface similarity was the best predictor of memory access and structural similarity was the best predictor of subjective soundness. This dissociation held not only between subjects, but also within subjects. That is, subjects given the soundness task immediately after the cued retrieval task judged that the very matches that had come to their minds most easily (the surface matches) were highly unsound (i.e., unlikely to be useful in inference). This suggests that similarity-based access may be based on qualitatively distinct processes from analogical inferencing.

It is not the case that higher-order relations contribute nothing to retrieval. Adding higher-order relations led to nonsignificantly more retrieval in two studies and to a small but significant benefit in the third. Other research has shown positive effects of higher-order relational matches on retrieval, especially in cases where subjects have been brought to do intensive encoding of the original materials (Faries & Reiser, 1988) or are expert in the domain (Novick, 1988a,b). But higher-order commonalities have a much bigger effect on mapping once the two analogs are present than they do on similarity-based retrieval, and the reverse is true for surface commonalities.

These results place several constraints on a computational model similarity-based retrieval. The first two criteria ensure that the model can provide an account of mapping and inference.

Structured representation criterion: The model must be able to store structured representations.
Structured mappings criterion: The model must incorporate processes of structural mapping (i.e., alignment and transfer) over its representations.

The remaining four criteria summarize the pattern of retrieval results:

Primacy of the mundane criterion: The majority of retrievals should be literal similarity matches: i.e., matches high in both structural and surface commonlaties.
Surface superiority criterion: Retrievals based on surface similarity are frequent.
Rare insights criterion: Relational remindings must occur at least occasionally, with lower frequency than literal similarity or surface remindings.
Scalability criterion: The model must be plausibly capable of being extended to large memory sizes.

No current model of transfer succeeds in satisfying all six criteria. There are two major approaches to memory models: Indexing models, commonly used in case-based reasoning work, and feature-vector models, commonly used in mathematical modeling of human memory. We examine the trade-offs of each in turn.

Most case-based reasoning models (Kass, 1986, 1989; Kolodner, 1984, 1988, 1989, 1993; Schank, 1982) use structured representations, and focus on the process of adapting and applying old cases to new situations. Such models satisfy the structured representation and structured mappings criteria. However, such models also typically presume a highly indexed memory in which the vocabulary used for indexing captures significant higher-order abstractions such as themes and principles. Viewed as psychological accounts, these models would predict that people should typically access the best structural match. That prediction fails to match the pattern of psychological results summarized by the primacy of the mundane and surface superiority criteria. Scalability is also an open question at this time, since no one has yet accumulated and indexed a large (1,000 to 10^6) corpus of structured representations.

The reverse set of advantages and disadvantages holds for approaches that model similarity as the result of a dot product (or some other simple operation) over feature vectors, as in many mathematical models of human memory (e.g., Gillund & Shiffrin, 1984; Hintzman, 1986, 1988; Medin & Schaffer, 1978; but see Murphy & Medin, 1985) as well as in many connectionist models of learning (e.g., Smolensky, 1988, see also the review by Humphreys, Bain & Pike (1989) and Ratcliff (1990).). These models typically utlize nonstructured knowledge representations and relatively simple match processes and hence do not allow for structural matching and inference. Such models also tend to utilize a unitary notion of similarity, an assumption that is called into question by the disassociation described above (see also Medin, Goldstone &

Gentner, 1993; Markman & Gentner, 1993). However, the use of feature-vectors has some advantages for modeling access to long-term memory. The computations are simple enough to make it feasible to compute many matches and choose the best, thus satisfying the scalability criterion. Further, because object features are included in the feature vectors, these models should be able to capture the surface superiority criterion and in many cases the primacy of the mundane criterion. (Failures on the latter will occur for cross-mappings, when the objects and relations match but their bindings do not.) It should be noted that some case-based reasoning work also restricts itself to feature-vector representations and thus has the same strengths and weaknesses (e.g., Stanfill & Waltz, 1986).

The MAC/FAC model seeks to combine the advantages of both approaches. We turn now to its description.

3. The MAC/FAC model

The complexity of the phenomena in similarity-based access suggests a two-stage model. Consider the computational constraints on access. The large number of cases in memory and the speed of human access suggests a computationally cheap process. But the requirement of judging soundness, essential to establishing whether a match can yield useful results, suggests an expensive match process. A common computational solution to such problems is to use a two-stage process, in which a cheap filter is used to pick out a subset of likely candidates for more expensive processing (c.f. Bareiss & King, 1989). MAC/FAC uses this strategy. The disassociation noted previously can be understood in terms of the interactions of its two stages.

Figure 1 illustrates the components of the MAC/FAC model. The inputs are a pool of memory items and a *probe*, i.e., a description for which a match is to be found. The output is an item from memory (i.e., a structured description) and a comparison of this item with the probe. (Section 3.1 below describes exactly what a comparison is.) Internally there are two stages. The MAC stage provides a cheap but non-structural filter, which only passes on a handful of items. The FAC stage uses a more expensive but more accurate structural match, to select the most similar item(s) from the MAC output and producing a full structural alignment. Each stage consists of *matchers*, which are applied to every input description, and a *selector*, which uses the evaluation of the matchers to select which comparisons are produced as the output of that stage. Conceptually, matchers are applied in parallel within each stage.

We make minimal assumptions concerning the global structure of long-term memory. We assume here only that there is a large pool of descriptions from which we must select one or a few that are most similar to a probe. We are uncommitted as to whether the pool is the whole of long-term memory or a subset selected via some other method, e.g., spreading activation.

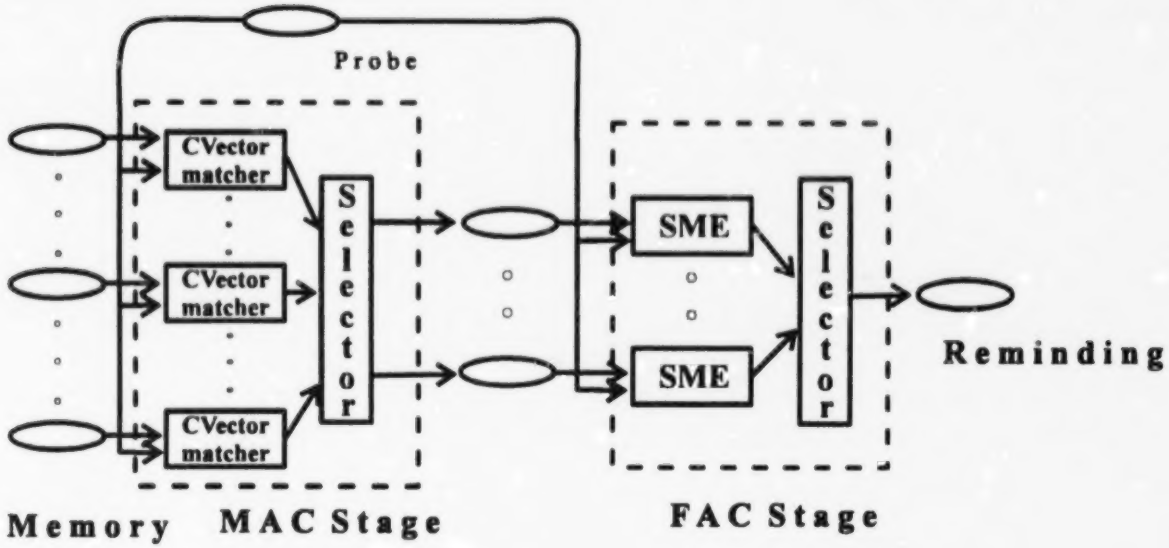

Figure 1: The MAC/FAC model

We begin by describing the FAC stage. In doing so, we also describe the computational framework which underlies MAC and FAC, including our conventions for representation and the information about the SME algorithm that is required to fully understand MAC/FAC.

3.1 The FAC stage and SME

The FAC stage takes as input the descriptions selected by the MAC stage, and computes a full structural match between each item and the probe. We model the FAC stage by using SME, the *Structure-Mapping Engine* (Falkenhainer, Forbus & Gentner, 1986, 1989). Here we briefly summarize SME's operation, both by way of describing the FAC stage and to provide the vocabulary needed to motivate the design of the MAC stage.

SME is an analogical matcher designed as a simulation of structure-mapping theory. It takes two inputs, a base description and a target description. (For simplicity we speak of these descriptions as being made up of items, meaning both objects and statements about these objects.) It computes a set of *global interpretations* of the comparison between base and target. Each global interpretation includes

- A set of *correspondences* which pair specific items in the base representation to specific items in the target.
- A *structural evaluation* reflecting the estimated soundness of the match. In subsequent processing, the structural evaluation provides one source of information about how seriously to take the match.
- A set of *candidate inferences*, potential new knowledge about the target which are suggested by the correspondences between the base and target. Candidate inferences are what give analogy its generative power, since they represent the importation of new knowledge into the target description. However, they are only conjectures; they must be tested and evaluated by other means.

We can illustrate these ideas with the Rutherford analogy, which describes the structure of the atom in terms of that of the solar system. The solar system is the base description and the atom is the target description.

- Rutherford paired the Sun to the nucleus and the planets to the electrons. These correspondences seem reasonable not because of intrinsic object similarities but because they allow various relational statements also to be placed in correspondence (i.e., *aligned*): e.g., the relative masses of the objects and the fact that the planets/electrons revolve around the Sun/nucleus.
- This interpretation is a selection from among many common relations. It focuses on the causal system of a central gravitational/electromagentic force, the relative mass of the two bodies within each system, and the fact that the less massive body revolves around the heavier body. Other common relations -- such as the relative temperatures or differences in color of the two objects -- that do not belong to a common connected system are not included in the interpretation. We refer to this preference for connected systems of common predicates as the *systematicity* principle.
- The preferred interpretation might also sanction new conjectures about the atom, such as that the cause of the electrons revolving around the nucleus is the existence of an attractive force.[3]

The interpretations produced by SME are structurally consistent, in that they satisfy the constraints of one-to-one mapping and parallel connectivity, as defined in Section 2.1. These constraints are important because they allow for the generation of coherent candidate inferences. The systematicity constraint is important because it captures the human preference for aligning connected systems of predicates (e.g., logical arguments or causal sequences). In addition, SME attempts to find *maximal* interpretations. An interpretation is maximal if adding any additional correspondences would render it structurally inconsistent. Maximality is important both because it reduces the number of possible interpretations and because it ensures that the full structural implications of a set of correspondences will be considered.

Before describing the SME algorithm further, some conventions concerning representation are in order. We use infix notation or Lisp prefix syntax for statements as appropriate. We use the term *functor* of a statement as a general term for the relation or function or connective that takes the remaining parts of the statement as its arguments. For example,

1. In `GREATER-THAN(HEIGHT(A),HEIGHT(B))`, the functor is `GREATER-THAN`.
2. In `NOT(ABOVE(B,A))`, the functor is `NOT`.
3. In `HEIGHT(A)`, the functor is `HEIGHT`.
4. In `RED(A)`, the functor is `RED`.

Example # 1 is an example of a *relation*. Relations range over truth values, and their arguments can be entities or other statements. Relations always have multiple arguments, with the exception of logical connectives (e.g., Example # 2), which are always treated as relations regardless of the number of arguments. For the purposes of structure-mapping, modal operators and other higher-order predicates are classifed as relations. Example # 3 is an example of a *function*, which maps one or more entities into another entity or constant. In our psychological modeling, functions are often used to represent known dimensions or components of structured objects (e.g., height, pressure, or color). Example # 4 is an

[3] Incorrect candidate inferences are also possible -- e.g. that the attractive force in the atom is gravity. What counts as a candidate inference versus as alignable (or non-alignable) structure depends on the reasoner's state of knowledge about the target.

example of an *attribute*, an atomic description of some property of an entity. Attributes take only one argument, to capture the notion of a unitary description. This of course does not mean that attributes cannot be decomposed. For instance, the following forms are logically equivalent:

- `RED(A)`
- `COLOR-OF(A, red)`
- `COLOR(A) = red`

However, we use these three distinct forms to represent distinct psychological constructs. Roughly, the first, an attribute, indicates that the subject thinks of redness as a quality of the object. The second, a relation, indicates that the subject has to some degree disengaged redness from the object and sees color as a relationship between an object and a set of possible values. The third, a function, indicates that the subject conceives of color as a dimension of general application and thinks of the color of A as a value along this dimension. We view this kind of dimensional representation as important because dimensions may in the process of comparison be aligned with quite different dimensions (e.g., `HEIGHT` and `DARKNESS`). Thus qualities that are conceived of as dimensions are more likely to participate in systematic cross-dimensional matches. (For the implications of this idea in analogical development, see Gentner & Rattermann, 1990; Gentner, Rattermann, Markman & Kotovsky, in press; Kotovsky & Gentner, 1988.)

With these conventions in mind, let us turn to the SME algorithm. SME operates via a local-to-global process. Conceptually, its operation can be divided into four phases. The first phase constructs a network of local matches between items in the base and target. The second phase constructs global interpretations by coalescing structurally consistent combinations of local matches. The third phase computes the structural evaluation, and the fourth phase computes candidate inferences for each interpretation. We examine each in turn.

SME begins by finding all possible local matches between statements in the base and statements in the target. A local match is created between base item B_i and target item T_j when either:

1. B_i and T_j are both statements whose functors are sufficiently alike (typically identical, but see below).
2. B_i and T_j are corresponding arguments of other statements which are connected by a local match and are both either objects or functions.

For instance, given the base item `B1` and the target item `T1` defined as:

B1: `(CAUSE Event17 Event31)`
T1: `(CAUSE Event5 Event63)`

a match would be hypothesized between `B1` and `T1` because their functors (i.e., `CAUSE`) are identical. This local match suggests in turn hypothesizing that `Event17` and `Event5` match, and also that `Event31` and `Event63` match. Each suggested match leads to the creation of new local matches involving the arguments of the statement if either (a) both are entities (e.g., objects or constants), (b) both are terms involving functions, which are an indirect means of referring to entities or dimensions, or (c) both are expressions whose functors match. Here is an example of substitution involving functions:

B2: `(PRESSURE Water32)`
T2: `(TEMPERATURE Brick45)`

B2 and T2 could be placed into correspondence if they were the arguments of some other matching pair of statements since PRESSURE and TEMPERATURE are both functions (in this case referring to values on physical dimensions of the respective objects).

The idea that two statements can match only if their relational predicates are "sufficiently alike" is based on the claim that some common relational content is required in analogy. We disagree with Holyoak and Thagard's (1989) claim that pure structural isomorphisms can qualify as analogies. They present the following pair:

Bill is smart and tall.	Rover is hungry and friendly.
Steve is smart.	Fido is hungry.
Tom is timid and tall.	Blackie is frisky and friendly.

Holyoak and Thagard (1989, p. 343) note that ACME (and five out of the eight subjects tested) could match this pair and agree on the best attribute correspondences. But the fact that it can be solved is not decisive: we would suggest that it is taken as a logical puzzle to be solved for the best correspondences, not as an analogy. The trouble with accepting pure graph matches is that it leads to the claim that pairs like (1) and (2) below are analogies, which seems patently untrue:

(1) Fred loves New York.

(2) General Motors sells cars.

Note that it is the *relational* meaning that must be shared; (2) and (3) form an analogy but (2) and (4) do not:

(3) Fred peddles popsicles.

(4) General Motors heads the list.

The question is how to formalize this requirement of common relational content. Structure-mapping uses the idea of *tiered identicality*. The default criterion for "sufficiently alike" for predicates other than functions is that the predicates are identical. We call this the *simple identicality* criterion. Simple identicality of conceptual relations is an excellent first-pass criterion because it is computationally cheap. The notion of smiple identicality might suggest an inability to process any matches other than literal matches. This is not the case. First, we assume that input representations are canonical *conceptual* representations, not semi-verbal strings. Second, functions, which represent domain dimensions, can be matched non-identically if they are embedded in matching relational structure. This ability to align non-identical functions provides considerable flexibility. This is what allows SME to make cross-dimensional matches, as when we interpret "Sally is sharper than Bill" to mean that Sally is smarter than Bill. However, there are circumstances where criteria requiring more processing are worthwhile (e.g., when placing two items in correspondence would allow a larger, or very relevant, structure to be mapped, as in Falkenhainer's (1987, 1990) work). In these circumstances weaker criteria (in that they allow more items to match) that involve more processing are allowed. One such test is *minimal ascension* (Falkenhainer, 1987, 1990) which allows two items to be placed into correspondence if their predicates have close common superordinates. Another technique is *decomposition:* Two concepts that are similar but not identical (such as "bestow" and "bequeath") are decomposed into a canonical representation language so

that their similarity is expressed as a partial identity (here, roughly "give"). Decomposition is the simplest form of *re-representation* (Gentner, 1989; Gentner & Rattermann, 1991), where additional knowledge is used to reformulate a description in order to achieve a better match. In this paper we only use SME with the first-level identicality constraint. As Section 6 argues, this simple constraint seems to provide a better psychological account than more complex constraints do.

The process of using matches to propose lower matches is recursive, ending with entity matches. SME does not try matches between every pair of objects in base and target: It only hypothesizes object matches when there is some aspect of the relational structure which suggests that the objects might correspond. This leads to substantial efficiencies over purely bottom-up matchers, such as (Winston 1980).

The output of the first phase is a network of *match hypotheses*, each representing a local match between an item of the base and target. At this stage the network is incoherent. The set of correspondences taken as a whole is structurally inconsistent, often including N-to-one mappings. Further, this initial network may contain match hypotheses that are not *grounded* and so can never be part of any global interpretation. A match hypothesis is *grounded* if a recursive chain of correspondences from it through its arguments exists all the way down to entities. Only grounded match hypotheses can participate in global interpretations. Otherwise, global interpretations might include statements whose arguments did not match, which would violate the parallel connectivity constraint.

Looking at a simple example makes this process clearer. Figure 2 shows two drawings used in psychological experiments concerning analogy,[4] with a propositional representation of these pictures suitable for simulation shown in Figure 3. The right hand side of Figure 3 shows the propositions in standard logical format, while the left hand side contains an equivalent graphical representation which is useful for understanding the match process. Figure 4 illustrates the match hypotheses computed by SME for these descriptions.

[4] We thank Arthur Markman for the drawings of Figure 2 and the corresponding representations.

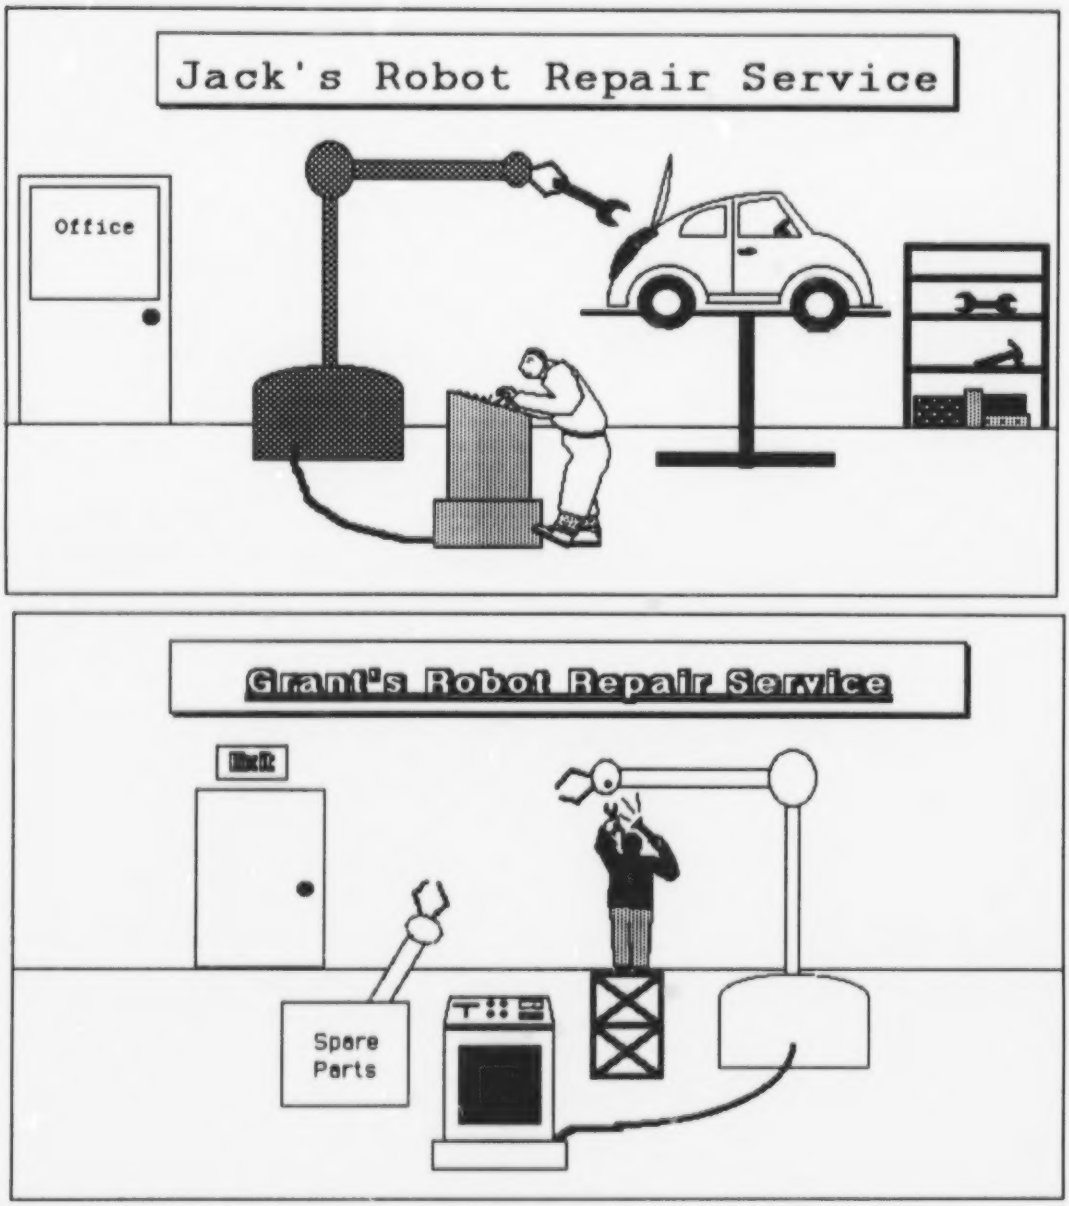

Figure 2: Two simple situations

Even though the initial network of match hypothesis is structurally inconsistent, it contains every consistent interpretation of the match; global interpretations emerge out of the initial network. Thus the maximum size of any global interpretation, as measured in number of correspondences, is limited by the size of this network. We exploit this fact in Section 3.3.

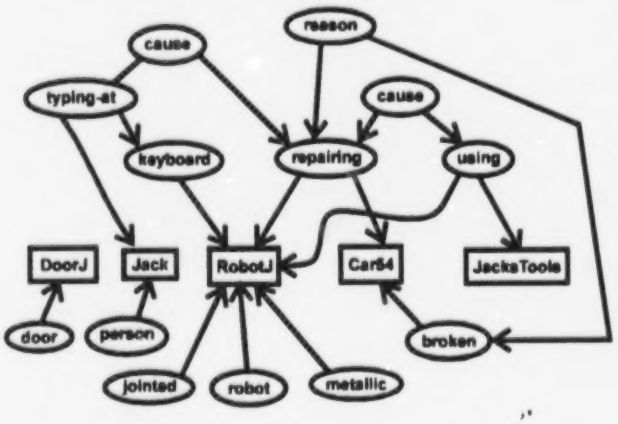

Jack's Robot Repair description:

(REASON (REPAIRING ROBOTJ CAR54)
(BROKEN CAR54))
(CAUSE (TYPING-AT JACK
(KEYBOARD ROBOTJ))
(REPAIRING ROBOTJ CAR54))
(CAUSE (REPAIRING ROBOTJ CAR54)
(USING ROBOTJ HANDTOOLSJ))
(DOOR DOORJ)
(JOINTED ROBOTJ)
(METALLIC ROBOTJ)
(ROBOT ROBOTJ)
(PERSON JACK)

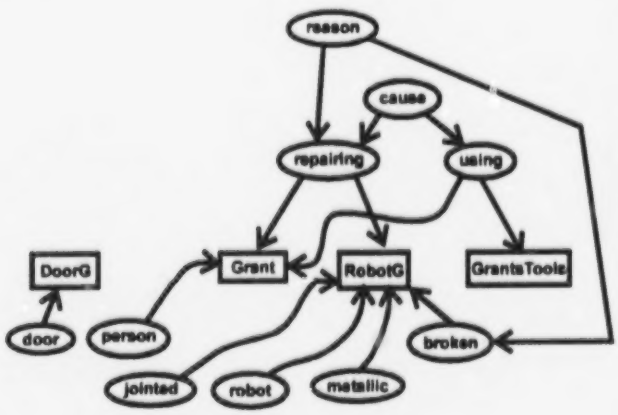

Grant's Robot Repair description:

(REASON (REPAIRING GRANT ROBOTG)
(BROKEN ROBOTG))
(CAUSE (REPAIRING GRANT ROBOTG)
(USING GRANT HANDTOOLSG))
(TOOL-SET HANDTOOLSG)
(DOOR DOORG)
(JOINTED ROBOTG)
(METALLIC ROBOTG)
(ROBOT ROBOTG)
(PERSON GRANT)

Here are two predicate calculus descriptions given to SME to illustrate the algorithm's operation.

Figure 3: Sample Descriptions

In the second phase, these local matches are coalesced into global interpretations. The SME algorithm combines structurally consistent combinations of match hypotheses (i.e., sets with consistent object bindings and consistent relational argument assignments). For instance, in Figure 4 there are two match hypotheses involving `Grant`, one which places him in correspondence with `Jack` because `PERSON` is true of both of them, and another match hypothesis which places `Grant` in correspondence with `RobotJ`, because both are agents of the same kind of action, `repairing`. No interpretation of this comparison can include both of these match hypotheses. Merging can be done exhaustively, producing all possible interpretations

(as in Falkenhainer, Forbus & Gentner 1986, 1989); however, we normally use a more psychologically plausible *greedy merge* algorithm, which produces only one or two interpretations and operates in linear time (Forbus & Oblinger 1990; Forbus, Ferguson, & Gentner, 1994).

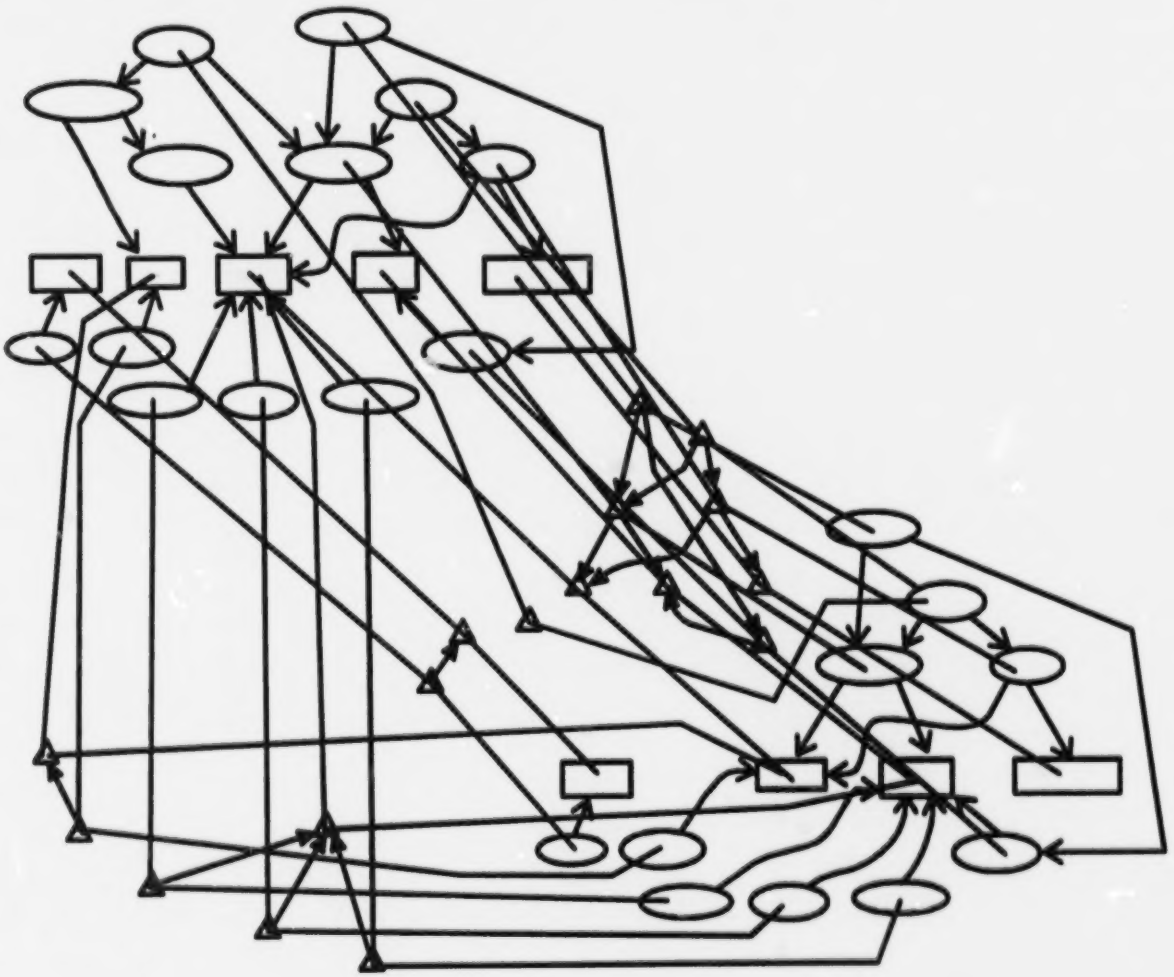

This picture illustrates the match hypotheses generated for a pair of simple descriptions. Match hypotheses are shown as triangles. Dashed lines indicate the base and target items each match hypothesis places in correspondence. The solid arrows leaving a match hypothesis indicates what others it relies upon to be structurally consistent. Notice that the one of the match hypotheses involving the occurrence of CAUSE in the target is structurally inconsistent, because its arguments cannot be aligned.

Figure 4: A match hypothesis forest

The third phase is structural evaluation. For simplicity, we describe this stage as conceptually distinct from the previous stage, although it is actually interleaved with building interpretations, since its results guide the greedy merge algorithm. To capture human preferences, the structural evaluation computation

should favor interpretations with many matches over those with few matches and deep interpretations over shallow interpretations. The first step is to assign an initial score to every match hypothesis. This helps enforce the size preference. The systematicity preference is implemented via a *trickle-down* method: match hypothesis scores are passed down to increment the scores of matching arguments.[5] That is, if $W(MH_1)$ is the score associated with a match hypothesis MH_1, MH_2 is a match hypothesis that applies to one of MH_1's arguments, and δ is the trickle-down factor, then $W(MH_2)$ is incremented as follows:

$$W(MH_2) \leftarrow \max(W(MH_2) + \delta W(MH_1);\ 1.0)$$

This local computation causes scores to cascade downwards, providing higher values to those object correspondences which support the alignment of large relational structures.. The structural evaluation of a global interpretation is simply the sum of the scores of the match hypotheses which comprise its correspondences.

```
GM1: 10 correspondences, SES = 4.66
  Object mappings:
   DOORG <-> DOORJ
   ROBOTG <-> CAR54
   GRANT <-> ROBOTJ
   HANDTOOLSG <-> HANDTOOLSJ
  No candidate inferences.
```

Here is a summary of the best interpretation for this match found by SME. `SES` refers to the structural evaluation of the interpretation.

Figure 5: Global interpretations for the example

The final phase is the computation of candidate inferences. Computing candidate inferences requires knowing the set of correspondences, so this takes place after the merge operation. Candidate inferences are generated by finding non-corresponding relational structure in the base which can be conjectured to hold in the target. The global interpretations built for the comparison of Figure 3 are shown in Figure 5. In this simple example there are no candidate inferences.

It is important to note that the literal similarity computation can produce purely relational interpretations as well as overall similarity interpretations, and that it can also produce purely surface interpretations as well. It is simply a question of which collection of local matches win. This reflects the human ability to process a novel comparison and discover only after the fact that it is an analogy. We assume that this all-purpose literal similarity mode is the normal mode of similarity processing in the absence of specific instructions. Consequently, SME creates initial local matches for attribute statements as well as for relational statements.

---

[5] The systematicity preference could have been implemented by differentially weighting matches at different levels. This method would seem to require a computationally implausible 'bird's eye' view of the representations. In a comparison of the two methods, the trickle-down method accounted for human soundness ratings better than treating weights directly as a function of order (Forbus & Gentner, 1989).

For SME to play a major role in a model of similarity-based retrieval, it should be consistent with psychological evidence. We have tested the psychological validity of SME as a simulation of analogical processing in several ways. For instance, we compared SME's structural evaluation scores with human soundness ratings for the Karla the hawk stories discussed below (Gentner & Landers, 1985; Rattermann & Gentner, 1987) Like human subjects, SME rated analogical matches higher than surface matches (Skorstad, Falkenhainer, and Gentner, 1987). The patterns of preference were similar across story sets: there was a significant positive correlation between the difference scores for SME and those for human subjects, where the difference score is the rating for analogy minus the rating for surface match within a given story set (Gentner, Rattermann & Forbus, 1993).

Since retrievals occur frequently, components in model of retrieval must be efficient. SME is quite efficient. The generation of match hypotheses is $O(n^2)$ on a serial machine, where n is the number of items in base or target, and should typically be better than O(log(n)) on data-parallel machines.[6] The generation of global interpretations is roughly O(log(n)) on a serial machine, using the greedy merge algorithm of Forbus & Oblinger (1990),[7] and even faster parallel merge algorithms seem feasible.

3.2 The FAC stage

The FAC stage is essentially a bank of SME matchers, all running in parallel in literal similarity mode.[8] These take as input the memory descriptions that are passed forward by the MAC stage and compute a structural alignment between each of these descriptions and the probe. The other component of the FAC stage is a *selector* -- currently a numerical threshold -- which chooses some subset of these comparisons to be available as the output of the retrieval system. (See Figure 1).

The FAC stage acts as a structural filter. It captures the human sensitivity to structural alignment and inferential potential (subject to the limited and possibly surface-heavy set of candidates provided by the MAC stage, as described below). Several remarks on this algorithm's role in retrieval are in order. We use the literal similarity algorithm, on the grounds that in reminding situations people can respond to and identify different kinds of similarity. (Recall that the literal similarity computation can compute relational similarity or object similarity as well as overall similarity). This choice seems ecologically sound because mundane matches are often reasonable guides to action; riding a new bicycle, for instance, is like riding other bicycles (Forbus & Gentner, 1986; Gentner, 1989; Medin & Ortony, 1989; Medin & Ross, 1989). Finally, this choice is necessary to model the high observed frequency of surface remindings. These surface remindings would mostly be rejected if FAC were strictly an analogy matcher. The selector for the FAC stage must choose a small set of matches for subsequent processing. Currently we select as output the best match, based on its structural evaluation, and any others within 10% of it. We settled on the 10% criteria because it generally returns a single result, only producing multiple results when there are two extremely close candidates. However, other criteria are possible and we have experimented with broadening the percentage, selecting a fixed number, selecting a maximum number (if capacity limits were assumed) and so forth. (One class of these experiments is described in Section 5.) We have also considered adding a threshold to the selector, so that if the best outcome is too weak the retrieval system returns nothing.

---

[6] The worst case parallel time would be O(n), in degenerate cases where all but one of the local matches is proposed by matching arguments.

[7] The original exhaustive merge algorithm was worst case factorial in the number of "clumps" of match hypotheses, but in practice was often quite efficient. See Falkenhainer, Forbus, & Gentner (1989) for details.

[8] In our current implementation SME is run sequentially on each candidate item in turn, but this is an artifact of the implementation.

3.3 The MAC stage

The MAC stage collects the initial set of matches between the probe and memory. Like the FAC stage, the MAC stage conceptually consists of a set of matchers and a selector that simply returns all items whose MAC score is within 10 % of the best score given that probe. The challenge of the MAC stage is in the design of its matcher. It must allow quickly comparing, in parallel, the probe to a large pool of descriptions and passing only a few on to the more expensive FAC stage. The rest of this section describes the design and implementation of the MAC matcher.

Let us start by examining in more detail the design criteria the MAC matcher must satisfy. Ideally, we would like the most similar or apt memory item for the given probe. Clearly running SME on the probe and every item in memory would provide the most accurate result. Unfortunately, even though SME is very efficient, it isn't efficient enough. SME operates by building intermediate structure, in the form of the network of local matches. The idea of building such networks for a pair of items, or a small number of pairs of items, is psychologically plausible, because the size of the match hypothesis network is polynomial in the size of the descriptions being matched. This means, depending on one's implementation assumptions, that a fixed-size piece of hardware could be built which could be dynamically reconfigured to represent any local match network for input descriptions of some bounded size. What is not plausible is that such networks could be built between a probe and every item in a large memory pool, and especially that this could happen quickly enough in neural architectures to account for observed retrieval times.

This architectural argument suggests that, while SME in literal similarity mode is fine for FAC, MAC must be made of simpler stuff. To escape having to suffer the complexity of the most accurate matcher in the "innermost loop" of retrieval, we must trade accuracy for efficiency. The MAC matcher must provide a crude, computationally cheap match process to pare down the vast set of memory items into a small set of candidates for more expensive processing. Ideally, MAC's computations should be simple enough to admit plausible parallel and/or connectionist implementations for large-scale memory pools.

What is the appropriate crude estimator of similarity? The most straightforward method would be to count the number of match hypotheses that FAC would generate in comparing a probe to a memory item. Let us call this number the *numerosity* of a comparison. Numerosity bears a rough relation to the potential size of the global interpretation, since the more local matches there are, the larger the global interpretation could potentially be. However, a large number of match hypotheses does not guarantee a large global interpretation, for two reasons. First, many match hypotheses might be ungrounded (recall Section 3.1) and hence cannot be part of any global interpretation. Second, often many combinations of match hypotheses are ruled out by the 1:1 constraint, preventing the formation of large global interpretations. Both reasons follow directly from the fact that numerosity is not structurally sensitive. However, something like numerosity is at least a crude estimate of similarity.

One straightforward way to implement a rough similarity estimator would be to calculate numerosity by building the actual match hypothesis network (e.g., to carry out the first part of a full analogy process) for the probe and each memory item, and then count the match hypotheses. This is what our original version of MAC/FAC did (Gentner, 1989b). It also is roughly what ARCS (Thagard et al 1990) does. ARCS models retrieval by building a network of connections similar to SME's match hypothesis network between the probe and each item in the memory pool that shares a semantically similar predicate with it.[9] As just discussed, we view these solutions as psychologically and computationally implausible. Even with parallel

[9] ARCS is based on Holyoak & Thagard's (1989) ACME, an analogy matcher which uses a localist connectionist network similar to SME's match hypothesis network to construct a single interpretation of a comparison via constraint satisfaction.

and/or neural hardware, it is hard to see how to generate match hypothesis networks between a probe and everything in a large pool of memory, while still providing realistic response times. A cheaper method is required.

We have developed a novel technique for estimating the degree of match in which structured representations are encoded as *content vectors*. Content vectors are flat summaries of the knowledge contained in complex relational structures. The content vector for a given description specifies which functors (i.e., relations, connectives, object attributes, functions, etc.) were used in that description and the number of times they occurred. Content vectors are assumed to arise automatically from structured representations and to remain associated with them. Content vectors are a special form of feature vectors.

More precisely, let Π be the set of functors used in the descriptions that constitute memory items and probes. We define the *content vector* of a structured description as follows. A content vector is an n-tuple of numbers, each component corresponding to a particular element of Π. Given a description ϕ, the value of each component of its content vector indicates how many times the corresponding element of Π occurs in ϕ. Components corresponding to elements of Π which do not appear in statements of ϕ have the value zero. One simple algorithm for computing content vectors is to simply to count the number of occurrences of each functor in the description. Thus if there were four occurrences of IMPLIES in a story, the value for the IMPLIES component of its content vector would be four. (Figure 6 illustrates.) Thus content vectors are easy to compute from a structured representation and can be stored economically (using sparse encoding, for instance, on serial machines).

Solar System: Structured representation

```
(CAUSE
  (GRAVITY (MASS SUN) (MASS PLANET))
  (ATTRACTS SUN PLANET))
(GREATER (TEMPERATURE SUN)
         (TEMPERATURE PLANET))
(CAUSE (AND (GREATER (MASS SUN)
                     (MASS PLANET))
            (ATTRACTS SUN PLANET))
       (REVOLVE-AROUND PLANET SUN))
```

Rutherford Atom: Structured representation

```
(CAUSE (OPPOSITE-SIGN (CHARGE NUCLEUS)
                      (CHARGE ELECTRON))
       (ATTRACTS NUCLEUS ELECTRON))
(REVOLVE-AROUND ELECTRON
                NUCLEUS)
(GREATER (MASS NUCLEUS)
         (MASS ELECTRON))
```

Solar System: Content Vector

```
(AND . 1)
(ATTRACTS . 1)
(CAUSE . 2)
(GRAVITY . 1)
(GREATER . 2)
(MASS . 2)
(OBJECTS . 2)
(REVOLVE-AROUND . 1)
(TEMPERATURE . 2)
```

Rutherford Atom: Content Vector

```
(ATTRACTS . 1)
(CAUSE . 1)
(CHARGE . 2)
(GREATER . 1)
(MASS . 2)
(OBJECTS . 2)
(OPPOSITE-SIGN . 1)
(REVOLVE-AROUND . 1)
```

Here are some simple predicate calculus representations and the corresponding content vectors. A simple counting algorithm is used here, in the simulation these are normalized to unit vectors.

Figure 6: Sample representations with content vectors

BEST COPY AVAILABLE

How good an approximation is the content vector dot product to what SME would produce? Suppose content vectors were generated using the simple counting algorithm described above. Then the product of each corresponding component is an overestimate of the number of match hypotheses that would be created between functors of that type, because it does not take into account the cases when the arguments to the match hypotheses could not be aligned. There is also a possibility of underestimation, since the dot product does not take into account matches between non-identical functions and entities, since discovering those matches requires tracing predicate bindings. However, in practice, the number of entity and non-identical function matches tends to be smaller than the number of ungrounded matches, so overall the dot product tends to overestimate numerosity, and hence will tend to be an overestimate of what SME would produce.

The dot product of content vectors provides exactly the computational basis the MAC stage needs. It could be implemented efficiently for large memories using a variety of massively parallel computation schemes. For instance, connectionist memories can be built which find the closest feature vector to a probe (Hinton & Anderson, 1989). Therefore the MAC stage can scale up.

To summarize, the MAC matcher works as follows: Each memory item has a content vector stored with it.[10] When a probe enters, its content vector is computed. A score is computed for each item in the memory pool by taking the dot product of its content vector with the probe's content vector. The MAC selector then produces as output the best match and everything within 10% of it, as described above. (As for the FAC stage, variants that could be considered include adding a bound on the number of items returned (to model capacity limitations) and implementing a threshold on the MAC selector so that if every match is too low MAC returns nothing).

Like other feature-vector schemes, the dot product of content vectors does not take the actual relational structure into account. It only calculates a numerical score, and hence doesn't produce the correspondences and candidate inferences which provide the power of analogical reasoning and learning. But the output of MAC feeds to the FAC stage, which operates on structured representations. Thus it is the FAC stage which both filters out structurally unsound remindings and produces the desired correspondences and candidate inferences. We claim that the interplay of the cheap but dumb computations of the MAC stage and the more expensive but structurally sensitive computations of the FAC stage explains the psychological phenomena of Section 2. As the first step in supporting this claim, we next demonstrate that MAC/FAC's behavior provides a good approximation of psychological data.

4. Cognitive Simulation Experiments

In this section we compare the performance of MAC/FAC with that of human subjects, using the "Karla the Hawk" stories (Gentner, Rattermann & Forbus, 1993; Rattermann & Gentner, 1987, Experiment 2). For these studies, we wrote sets of stories consisting of base stories plus four variants, created by systematically varying the kind of commonalities. All stories shared first-order relations (primarily events), but varied in which other commonalities were present, as shown in Table 1. The LS (literal similarity) stories shared both higher-order relational structure and object attributes. The AN stories shared higher-order relational structure but contained different attributes, while the SF stories shared attributes but

[10] We normalize content vectors to unit vectors, both to reduce the sensitivity to overall size of the descriptions and because we assume that psychologically plausible implementation substrates for MAC/FAC (e.g., neural systems) will involve processing units of limited dynamic range.

contained different higher-order relational structure. The FOR stories differed both in attributes and higher-order relational structure.

	Common 1st o. relations	Common h.o. relations	Common object attributes
LS:	Yes	Yes	Yes
SF:	Yes	No	Yes
AN:	Yes	Yes	No
FOR:	Yes	No	No

Table 1: Types of stories used in the "Karla the Hawk" experiments

In this study, the subjects were first given 32 stories to remember, of which 20 were base stories and 12 were distractors. They were later presented with 20 probe stories which matched the base stories as follows: five LS matches, five AN matches, five SF matches and five FOR matches. They were told to write down any prior stories of which they were reminded. (Which stories were in each similarity condition was varied across subjects.) As shown in Table 2 the proportions of remindings for different match types were .56 for LS, .53 for SF, .12 for AN and.09 for FOR. Table 2 also shows that this retrievability order has been stable acorss three variations of this study: LS ≥ SF > AN ≥ FOR.[11]

Condition	Proportion
LS	0.56
SF	0.53
AN	0.12
FOR	0.09

Table 2: Proportion of remindings for different match types: Human subjects

As discussed above, this retrievability order differs strikingly from the soundness ordering. When subjects were asked to rate how *sound* the matches were -- how well the inferences from one story would apply to the other -- they rated analogy (AN) and literal similarity (LS) as significantly more sound than surface similarity (SF) and FOR matches (matches based only on common first-order relations, primarily events). SME running in analogy mode on SF and AN matches correctly reflected human soundness rankings (Forbus & Gentner, 1989; Gentner et al, in press; Skorstad et al, 1987). Here we seek to capture human retrieval patterns: Does MAC/FAC duplicate the human propensity for retrieving SF and LS matches rather than AN and FOR matches. The idea is to give MAC/FAC a memory set of stories, then probe with various new stories. To count as a retrieval, a story must make it through both MAC and FAC. We use replication of the ordering found in the psychological data, rather than the exact percentages, as our

[11] LS and SF did not differ significantly in retrievability. In Experiment 2, AN and FOR did not differ significantly, although in Experiment 1 AN matches were better retrieved than FOR matches.

criterion for success because this measure is more robust, being less sensitive to the detailed properties of the databases.

For the computational experiments, we encoded predicate calculus representations for 9 of the 20 story sets (45 stories). Figure 7 shows one of the story representations. These stories are used in all three experiments described below.

Figure 7: A representation from the Karla The Hawk story set

```
(FOLLOW
  (PROMISE MAN1 KARLA
    (NOT (ATTACK MAN1 KARLA)))
  (ATTACK MAN1 DEER))
(CAUSE
 (EQUALS (HAPPINESS MAN1) HIGH)
 (PROMISE MAN1 KARLA
   (NOT (ATTACK MAN1 KARLA))))
(CAUSE
 (OBTAIN MAN1 FEATHERS)
 (EQUALS (HAPPINESS MAN1) HIGH))
(FOLLOW
 (OFFER KARLA FEATHERS MAN1)
 (OBTAIN MAN1 FEATHERS))
(CAUSE
 (REALIZE KARLA
     (DESIRE MAN1 FEATHERS))
 (OFFER KARLA FEATHERS MAN1))
(FOLLOW
 (EQUALS
  (SUCCESS
     (ATTACK MAN1 KARLA)) F)
 (REALIZE KARLA
   (DESIRE MAN1 FEATHERS)))
(CAUSE
 (NOT (USED-FOR
        FEATHERS CROSS-BOW))
 (EQUALS (SUCCESS
           ATTACK MAN1 KARLA))
               F))
(FOLLOW
 (ATTACK MAN1 KARLA)
 (EQUALS (SUCCESS
           (ATTACK MAN1 KARLA))
         F))
(FOLLOW
  (SEE KARLA MAN1)
  (ATTACK MAN1 KARLA))
(HAPPEN (SEE KARLA MAN1))
(LIVES KARLA LOC1)
(POSSESS MAN1 CROSS-BOW)
(POSSESS KARLA FEATHERS)
(RUMINANT DEER)
(ANTLERED DEER)
(HOOFED DEER)
(QUADRIPED DEER)
(MAMMAL DEER)
(THIN CROSS-BOW)
(LARGE CROSS-BOW)
(MEDIEVAL CROSS-BOW)
(WOODEN CROSS-BOW)
(WEAPON CROSS-BOW)
(BLACK FEATHERS)
(CO'/ERING FEATHERS)
(LONG FEATHERS)
(SOFT FEATHERS)
(ASSET FEATHERS)
(VOCAL MAN1)
(BIPED MAN1)
(HUNTER MAN1)
(WARLIKE MAN1)
(HUMAN MAN1)
(MALE MAN1)
(PREDATORY KARLA)
(BLACK KARLA)
(POWERFUL KARLA)
(LARGE KARLA)
(HAWK KARLA)
```

4.1 Cognitive Simulation Experiment 1

In our first study, we put the 9 base stories in memory, along with the 9 FOR stories which served as distractors. We then used each of the variants -- LS, SF, and AN -- as probes. This roughly resembles the original task, but MAC/FAC's job is easier in that (1) it has only 18 stories in memory, while subjects had 32, in addition to their vast background knowledge; (2) subjects were tested after a week's delay, so that there could have been some degradation of the memory representations.

Table 3 shows the proportion of times the base story made it through MAC and (then) through FAC. The FAC output is what corresponds to human retrievals. MAC/FAC's performance is much better than that of the human subjects, perhaps partly because of the differences noted above. However, the key point is that its results show the same ordering as those of human subjects: LS > SF > AN.

Probes	MAC	FAC
LS	1.0	1.0
SF	0.89	0.89
AN	0.67	0.67

Memory contains 9 base stories and 9 FOR matches; probes were the 9 LS, 9 SF, and 9 AN stories.
The rows show proportion of times the correct base story was retrieved for different probe types.

Table 3: Proportion of correct retrievals given different kinds of probes

4.2 Cognitive Simulation Experiment 2

To give MAC/FAC a harder challenge, we put the four variants of each base story into memory. This made a larger memory set (36 stories) and also one with many competing similar choices. Each base story in turn was used as a probe. This is almost the reverse of the task subjects faced, and is more difficult.

Table 4 shows the mean number of matches of different similarity types that succeed in getting through MAC and (then) through FAC. There are several interesting points to note here. First, the retrieval results (i.e., the number that make it through both stages) ordinally match the results for human subjects: LS > SF > AN > FOR. This degree of fit is encouraging, given the difference in task. Second, as expected, MAC produces some matches that are rejected by FAC. This number depends partly on the criteria for the two stages. Here, with MAC and FAC both set at 10%, the mean number of memory items produced by MAC is 3.4, and the mean number accepted by FAC is 1.6. Third, as expected, FAC succeeds in acting as a structural filter on the MAC matches. It accepts all of the LS matches MAC proposes and some of the partial matches (i.e., SF and AN), while rejecting most of the inappropriate matches (i.e., FOR and matches with stories from other sets).

Retrievals	MAC	FAC
LS	0.78	0.78
SF	0.78	0.44
TA	0.33	0.22
FOR	0.22	0.0
Other	1.33	0.22

Memory contains 36 stories (LS, SF, AN, and FOR for 9 story sets); the 9 base stories used as probes
Other = any retrieval from a story set different from the one to which the base belongs.

Table 4: Mean numbers of different match types retrieved per probe when base stories are used as probes

4.3 Cognitive Simulation Experiment 3

In the prior simulations, LS matches were the resounding winner. While this is reassuring, it is also interesting to know which matches would be retrieved if there were no perfect overall matches. Therefore we removed the LS variants from memory and repeated the second simulation experiment, again probing

with the base stories. As Table 5 shows, SF matches are now the clear winners in both the MAC and FAC stages. Again, the ordinal results match well with those of subjects: SF > AN > FOR.

Retrievals	MAC	FAC
SF	0.88	0.78
AN	0.56	0.56
FOR	0.22	0.11
Other	1.11	0.11

Memory contains 27 stories (9 SF, 9 AN, 9 FOR); 9 base stories used as probes

Table 5: Mean numbers of different match types retrieved per probe with base stories as probes and no LS stories in memory

4.4 Summary of Cognitive Simulation Experiments

The results are encouraging. First, MAC/FAC's retrieval results (i.e., the number that make it through both stages) ordinally match the results for human subjects: LS > SF > AN > FOR. Second, as expected, MAC produces some matches that are rejected by FAC. The mean number of memory items produced by MAC is 3.4 and the mean number accepted by FAC is 1.6. Third, FAC succeeds in its job as a structural filter on the MAC matches. It accepts all of the LS matches proposed by MAC and some of the partial matches (the SF, AN and FOR matches), and rejects most of the inappropriate matches (the "other" matches from different story sets). It might seem puzzling that FAC accepts more SF matches than AN matches, when it normally would prefer AN over SF. The reason is that it is not generally being offered this choice. Rather, it must choose the best from the matches passed on by MAC for a given probe (which might be AN and LS, or SF and LS, for example).

It is useful to compare MAC/FAC 's performance with that of Thagard, Holyoak, Nelson & Gochfeld's (1990) ARCS model of similarity-based retrieval, the most comparable alternate model. Thagard et al gave ARCS the Karla the hawk story in memory along with 100 fables as distractors. When given the four similarity variants as probes, ARCS produced asymptotic activations as follows: LS (.67), FOR (-.11), SF (-.17), AN (-.27). ARCS thus exhibits at least two violations of the LS ≥ SF > AN ≥ FOR order found for human remindings. First, SF remindings, which should be about as likely as LS remindings, are quite infrequent in ARCS -- less frequent than even the FOR matches. Second, AN matches are less frequent than FOR matches in ARCS, whereas for humans AN was always ordinally greater than FOR and (in Experiment 1) significantly so. Thus MAC/FAC explains the data better than ARCS. This is especially interesting because Thagard et al argue that a complex localist connectionist network which integrates semantic, structural, and pragmatic constraints is required to model similarity-based reminding. While such models are intriguing, MAC/FAC shows that a simpler model can provide a better account of the data. We compare MAC/FAC with ARCS in more detail in Section 6.

Finally, and most importantly, MAC/FAC's overall pattern of behavior captures the motivating phenomena. It allows for structured representations and for processes of structural alignment and mapping over these representations thus satisfying the *structural representation* and *structured mappings* criteria. It produces fewer analogical matches than literal similarity or surface matches, thus satisfying the *existence of rare insights* criterion. The majority of its retrievals are LS matches, thus satisfying the *primacy of the*

mundane criterion. It also produces a fairly large number of SF matches, thus satisfying the *surface superiority* criterion. Finally, its algorithms are simple enough to apply over large-scale memories, thus satisfying the *scalability* criterion.

5. Sensitivity Analyses

The experiments of the previous section show that the MAC/FAC model can account for psychological retrieval data. This section looks more closely into *why* it does, by seeing how sensitive the results are to different factors in the model. These analyses are similar in spirit to those carried out by VanLehn (1989) in his SIERRA project. VanLehn used his model to generate different possible learning sequences to see if these variations covered the space of observed mistakes made by human learners in subtraction problems. Thus variations in the model were used to generate hypotheses about the space of individual differences. Our methodology is quite similar, in that we vary aspects of our model in order to better understand how it accounts for data. The key difference is that we are not attempting to model individual differences, but instead are investigating how our results depend on different aspects of the theory. Such *sensitivity analyses* are routinely used in other areas of science and engineering; we believe they are also an important tool for cognitive modeling.

Sensitivity analyses can provide insight into why a simulation works. Any working cognitive simulation rests on a large number of design choices. Examples of design choices include the setting of parameters, the kinds of data provided as input, and even the particular algorithms used. Some of these design choices are forced by the theory being tested, some choices are only weakly constrained by the theory, and others are irrelevant to the theory being tested, but are necessary to create a working artifact. Sensitivity analyses can help verify that the source of a simulation's performance rests with the theoretically important design choices. Varying theoretically forced choices should lead to a degradation of the simulation's ability to replicate human performance. Otherwise, the source of the performance lies elsewhere. On the other hand, varying theoretically irrelevant choices should not affect the results, and if they do, it suggests that something other than the motivating theory is responsible for the simulator's performance. Finally, seeing how the ability to match human performance varies with parameters that are only weakly constrained by theory can lead to insights about why the model works.

In the rest of this section, we describe a series of sensitivity experiments on MAC/FAC. These experiments demonstrate that its ability to replicate human performance is robust, and that this ability depends crucially on the theoretically important design choices. We first describe the methodology used in these experiments in detail, and then describe three sensitivity analyses.

5.1 Method for Sensitivity Analyses

A sensitivity analysis requires a standard of comparison, a baseline against which to judge the results of variations. We use as our baseline the simulation experiments described in Section 4. We say that a particular set of design choices *satisfies the data* if re-running the simulation experiments with that set of design choices yields results which match the human data. That is, the frequency of retrievals must follow the pattern LS ≥ SF > AN > FOR.

There are many design choices which could be explored via sensitivity analyses. Conceptually, one can think of sets of design choices as points in a high dimensional space. In essense, the simulation studies of Section 4 provide information about one point in the design space. This metaphor is excellent for choices of numerical parameters, since these dimensions can be viewed as continuous. This metaphor is not as

useful for other kinds of design choices, e.g., algorithmic choices, since systematically enumerating the set of plausible algorithms for a task is quite difficult. To best visualize the results, choosing two numerical dimensions to vary allows patterns of satisfaction to be displayed as a table, whose entries represent measurements of the ability of the model to satisfy the data at sampled points in the design space.

The two most interesting numerical parameters with respect to sensitivity analyses on MAC/FAC are the selector widths for the MAC and FAC stages, since these are only weakly constrained by the theory. They should be narrow, in order to reject inappropriate remindings, but we currently see no theoretically motivated method to calculate precise predictions for these parameters. Therefore in these analyses we use an empirical approach. We vary the selector widths, using these variations as the axes of a subset of the design space. Recall that a selector of width W accepts all matches within W % of the largest input. That is, a selector with width 10 % outputs the best match plus any other matches that are within 10 % of the score of the best match, while a selector of width 100 % will simply pass through all of its inputs. In the experiments below, selector widths for both MAC and FAC are varied from 1 to 100 %, in 10 % increments. Each entry in the table indicates whether that pair of width settings, combined with the other design choices, satisfied the data. When the pattern of retrieval is violated, the table entry contains information about the particular kind of violation.

Viewed as a map, the table of results from the sensitivity analysis can be divided into *viable regions*, subspaces of design choices which allow the model to satisfiy the data, and *non-viable regions*, where they do not. The existence of viable regions is of course critical for a successful simulation. However, the nature of the non-viable regions is also interesting, because they provide a source of insight into why the model works. Seeing how a bridge collapses after replacing a particular strut with a weaker material (preferably via simulation) supports the conclusion that the strength of that strut was a factor in preventing collapse.

It should be noted that the computational costs of these experiments is large but not horrendous. Essentially, the cognitive simulation experiments of Sections 4.1and 4.2 were replicated for each pair of selector widths, i.e., 121 times. Each repetition required running the MAC matcher 810 times,[12] for a total of 98,010 times. The number of FAC executions varies with the size of the set output from MAC, of course, and varied substantially according to the particular design choices made (as shown below). A reasonably accurate estimate for the lower bound of FAC executions for each experiment was 900 and a reasonable upper bound is 1,600. The MAC matcher takes roughly 0.002 seconds for each pair of content vectors and the FAC matcher (i.e., SME) takes between 1.0 and 11 seconds for each pair of structured representations, with an average time of roughly 4 seconds.[13] Thus the time to run MAC/FAC for each probe typically ranges from 3 to 10 seconds. A naive system for doing sensitivity analyses could use as much as five hours per analysis (14,000 seconds for MAC, 4,000 seconds for FAC). However, we found that by caching the results of matches in a simple database, we could cut the CPU requirements for these analyses considerably.

---

[12] The first experiment involves 486 MAC executions because there are 18 stories in memory and 27 probes. The second experiment involves 324 MAC executions because there are 36 stories in memory and 9 probes.

[13] These times are for an IBM RS/6000 Model 350, SME3b, which was used in all experiments in this section. An earlier version of SME was used in Forbus & Gentner 1991, and in the experiments in Section 6.

5.2 Sensitivity Analysis One: Robustness

In this experiment we tested the robustness of MAC/FAC's ability to satisfy the data by varying the selector widths. Table 6 shows the results. Notice that there is one region which satisfies the data: When the MAC width is between 10 % and 20 % and FAC is at least 10 %. The moderately large viable subspace indicates that MAC/FAC's performance is robust, and not hostage to a particular choice of selector width settings.

Rows: MAC widths. Columns: FAC widths.

	1%	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
1%	4	4	4	4	4	4	4	4	4	4	4
10%	4	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
20%	4	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
30%	4	64	64	80	80	80	112	112	112	112	112
40%	4	64	64	80	80	80	112	112	368	368	368
50%	4	64	64	80	80	80	112	368	368	368	368
60%	4	64	64	80	80	80	112	368	368	368	368
70%	4	64	64	80	80	80	112	368	368	368	368
80%	4	64	64	80	80	80	112	368	368	368	368
90%	4	64	64	80	80	80	112	368	368	368	368
100%	4	64	64	80	80	80	112	368	368	368	368

Legend
Y = Satisfies the data
4 = No analogies
64 = SF ≤ AN
80 = LS ≤ AN , SF ≤ AN
112 = LS ≤ AN , SF ≤ AN , LS ≤ FOR
368 = LS ≤ AN , SF ≤ AN , LS ≤ FOR , LS ≤ DT

Rows are the width of the MAC selector, Columns correspond to the width of the FAC selector. The codes describe whether that combination of selector widths allows MAC/FAC to account for the human data, and if not, what criteria were violated.

Table 6: Sensitivity to selector width, normalized content vectors

As discussed above, it is important to show that there are parameter settings that do not fit the human data, to establish that the theoretical variables actually matter. When either MAC or FAC is too narrow (i.e., MAC of 1 % or FAC of 1 %), analogies are never retrieved. This violates the rare insights criterion. When MAC is broad (30 % or larger), making FAC too broad leads first to too many analogies, and then to junk remindings. The shape of the region of viability suggests that while FAC is necessary to provide structural matching and candidate inferences, MAC provides most of the filtering. Since that is MAC's intended purpose, this provides further evidence that the simulation works according to the principles of its design, rather than some unknown factor.

The evidence that the results are not very sensitive to the particular choice of selector widths in the original experiments (i.e., 10 % for both MAC and FAC) is reassuring. The next two sensitivity analyses explore other design choices, using the same methodology as this experiment.

5.3 Sensitivity Analysis Two: Irrelevance of normalization details

In other sensitivity experiments on analogical processing algorithms (Forbus & Gentner, 1989), we demonstrated that the choice of normalization algorithm could affect outcomes in simulations of structural evaluation in comparisons. The purpose of this analysis is to determine if our design choice of using unit content vectors (see Section 3.3) was a significant factor in our results.

To explore this question, we consider two variations on the content vector representation. The first variation is simply not to use any kind of normalization at all. That is, we simply use as the strength of each component of the content vector the number of statements and terms which contained the corresponding predicate. (The computation of normalized content vectors involves an additional step, of dividing each component by the total number of predicates in the description.) The results of this manipulation are illustrated in Table 7. The key point to notice about this table is that the viable region is roughly the same size and shape as the viable region for normalized content vectors. This lends support to the claim that the outcome of the simulation experiments is not heavily determined by the particular normalization algorithm chosen.

Rows: MAC widths. Columns: FAC widths.

	1%	10 %	20 %	30 %	40 %	50 %	60 %	70 %	80 %	90 %	100 %
1%	4	4	4	4	4	4	4	4	4	4	4
10%	4	Y	Y	Y	Y	Y	Y	Y	Y	256	256
20%	4	64	Y	Y	Y	Y	Y	256	256	256	256
30%	4	64	64	336	336	336	336	336	336	336	336
40%	4	64	Y	80	88	88	120	376	376	376	376
50%	4	64	64	80	80	80	112	368	368	368	368
60%	4	64	64	30	80	80	112	368	368	368	368
70%	4	64	64	80	80	80	112	368	368	368	368
80%	4	64	64	80	80	80	112	368	368	368	368
90%	4	64	64	80	80	80	112	368	368	368	368
100 %	4	64	64	80	80	80	112	368	368	368	368

Legend
4 = No analogies
64 = SF ≤ AN
80 = LS ≤ AN, SF ≤ AN
88 = LS ≤ SF, LS ≤ AN , SF ≤ AN
112 = LS ≤ AN, LS ≤ FA, SF ≤ AN
120 = LS ≤ SF, LS ≤ AN, LS ≤ FA, SF ≤ AN
256 = LS ≤ DT
336 = LS ≤ AN, SF ≤ AN, LS ≤ DT
368 = LS ≤ AN, LS ≤ FA, SF ≤ AN, LS ≤ DT
376 = LS ≤ SF, LS ≤ AN, LS ≤ FA, SF ≤ AN, LS ≤ DT

Rows are the width of the MAC selector, Columns correspond to the width of the FAC selector. The codes describe whether that combination of selector widths allows MAC/FAC to account for the human data, and if not, what criteria were violated.

Table 7: Sensitivity analysis, unnormalized content vectors

The second variation we consider is to change what aspect of the overlap content vectors measure. Recall that the idea of content vectors is to compare the pattern of functors which appear in two structured descriptions. There are several ways to characterize such patterns. The MAC/FAC design choice, normalized content vectors, estimates the overlap in terms of the relative frequency of functors in the two descriptions, independent of their sizes. The unnormalized content vectors just examined estimate the total size of the overlap. But is it the pattern of overlap that is relevant, or just how many functors two descriptions have in common? We can investigate this question by changing the structure of content vectors so that they represent only the set of predicates that are used in the structured representation, without regard to number of occurrences. We call this variation *binary content vectors* since each component is essentially a one bit answer to the question of whether the structured represenation contains or does not contain a particular predicate. Thus the dot product of two binary content vectors is a measure of the overlap in number of shared predicates. (Again, we normalize to unit vectors, both to avoid size biases and because we assume that psychologically plausible implementation substrates (e.g., neural systems) will have limited dynamic range.) The results of this manipulation are shown in Table 8.

Rows: MAC widths. Columns: FAC widths.

	1%	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
1%	4	4	4	4	4	4	4	4	4	4	4
10%	4	4	4	4	4	4	4	4	4	4	4
20%	4	4	4	4	4	4	4	4	4	4	4
30%	4	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
40%	4	Y	Y	64	Y	Y	Y	Y	Y	Y	Y
50%	4	64	Y	64	Y	Y	Y	Y	Y	Y	Y
60%	4	64	64	80	80	80	112	112	112	112	112
70%	4	64	64	80	80	80	112	112	368	368	368
80%	4	64	64	80	80	80	112	368	368	368	368
90%	4	64	64	80	80	80	112	368	368	368	368
100%	4	64	64	80	80	80	112	368	368	368	368

Legend
Y = Syslit predictions satisfied
4 = No analogies
64 = SF ≤ AN
80 = LS ≤ AN, SF ≤ AN
112 = LS ≤ AN, LS ≤ FA, SF ≤ AN
368 = LS ≤ AN, LS ≤ FA, SF ≤ AN, LS ≤ DT

Binary content vectors measure the size of overlap in predicates. As before, rows are the width of the MAC selector, Columns correspond to the width of the FAC selector. The codes describe whether that combination of selector widths allows MAC/FAC to account for the human data, and if not, what criteria were violated.

Table 8: Sensitivity analysis for binary content vectors

Again, the overall pattern of results is the same: With selector widths that are too narrow no analogies are retrieved, and with selector widths that are too broad, too many analogies are retrieved, followed as widths increase by too many "junk" retrievals. The interesting difference is that the region for the selector widths has changed: The viable wide-FAC range lies with MAC between 30--50 %, whereas it was between 10--20 % for the original content vectors. Comparing the average number of representations output by MAC

for these ranges provides some insight as to why this should be so: For binary content vectors the average output size was 2, for standard content vectors the average was 1.5. In both cases, the next step of MAC selector width allows, on the average, another representation to make it through to FAC. Yet one more step in MAC selector width allows many more representations to get through to FAC. Thus measuring only the number of shared predicates shifts the viable region, but does not substantially change its character.

From these two analyses, we conclude that the choice of normalization algorithm does not substantively affect the results. Since the normalization algorithm is not a theoretically determined choice, these analyses support the conclusion that the simulation works according to the theoretical account.

5.4 Sensitivity Analysis Three: Attributes versus relations

Content vectors homogenize structured representations. They unify information about attributes of objects, relationships between objects, and argument structure. Is including every kind of information in content vectors necessary? Given the frequency of literal similarity and surface feature matches, both of which share many attributes, a possible hypothesis is that content vectors could be built using attributes alone. On the other extreme, the approachs used in case-based reasoning tend to ignore attributes, and use only relational information. To mimic these approaches in MAC/FAC, we could use content vectors which leave out attributes, and include only relational predicates. This analysis explores both of these extreme hypotheses.

To explore the degree to which using attribute information only in content vectors would allow MAC/FAC to satisfy the data, we modified the algorithm which computes content vectors to ignore anything other than attributes. The results of the sensitivty analysis are shown in Table 9. The pattern of results is dramatically different than in previous experiments. There is no viable region at all. This experiment provides strong evidence that using attribute information alone in content vectors cannot satisfy the data.

The failure of attributes alone to provide adequate filtering may not be surprising. Is relational information alone enough? To explore this question we again modified the algorithm that computes content vectors, this time to not include attributes. These new content vectors therefore only contained relationships between objects and higher-order relations, such as logical connectives. The same methodology for the sensitivity analysis was followed.

The results of the sensitivity analysis are shown in Table 10. Like the attribute-only content vectors, the relation-only content vectors also fail to satisfy the data in a psychologically plausible manner, but for different reasons. Almost uniformly, that is, when either the MAC width is less than 40 % or when the FAC width is greater than 70 %, more "junk" matches come through -- stories from other sets, and FOR stories (i.e., those which match only in terms of first-order relations and not attributes or causal structure). The region where the data is not satisfied and the MAC width ranges between 20 % and 70 % is very much like the failures that occur for the attribute-only vectors (e.g., more analogies retrieved with narrow FAC than psychologically plausible). There is in fact a region in Table 10 where the pattern of results matches the human data, when the MAC width is between 40% and 50% and the FAC width ranges from 20% to either 60% or 70%. However, the size of the MAC output in this range is roughly one half of the total size of the memory pool. Consequently, this is not a viable region, because it demands far too much of FAC.

Rows: MAC widths. Columns: FAC widths.

	1%	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
1%	326	322	322	322	322	322	322	322	322	322	322
10%	326	336	336	336	336	336	336	336	336	336	336
20%	260	320	320	336	336	336	336	336	336	336	336
30%	4	64	320	336	344	344	344	344	344	344	344
40%	4	64	64	336	336	336	368	368	368	368	368
50%	4	64	64	80	88	88	120	376	376	376	376
60%	4	64	64	80	80	80	112	368	368	368	368
70%	4	64	64	80	80	80	112	368	368	368	368
80%	4	64	64	80	80	80	112	368	368	368	368
90%	4	64	64	80	80	80	112	368	368	368	368
100%	4	64	64	80	80	80	112	368	368	368	368

Legend

Y = Syslit predictions satisfied
4 = No analogies
64 = SF ≤ AN
80 = LS ≤ AN, SF ≤ AN
88 = LS ≤ SF, LS ≤ AN, SF ≤ AN
112 = LS ≤ AN, LS ≤ FA, SF ≤ AN
120 = LS ≤ SF, LS ≤ AN, LS ≤ FA, SF ≤ AN
60 = No analogies, LS ≤ DT 320 = SF ≤ AN, LS ≤ DT
322 = No surface matches, SF ≤ AN, LS ≤ DT
326 = No surface matches, no analogies, SF ≤ AN, LS ≤ DT
336 = LS ≤ AN, SF ≤ AN, LS ≤ DT
344 = LS ≤ SF, LS ≤ AN, SF ≤ AN, LS ≤ DT
368 = LS ≤ AN, LS ≤ FA, SF ≤ AN, LS ≤ DT
376 = LS ≤ SF, LS ≤ AN, LS ≤ FA, SF ≤ AN, LS ≤ DT

These results obtained are using content vectors which only included attributes, leaving out relations and logical connectives. As before, rows are the width of the MAC selector, Columns correspond to the width of the FAC selector. The codes describe whether that pair of selector widths allows MAC/FAC to account for the human data, and if not, what criteria were violated.

Table 9 Sensitivity analysis of attribute-only content vectors

Rows: MAC widths. Columns: FAC widths.

	1%	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
1%	260	260	260	268	268	268	268	268	268	268	268
10%	270	268	268	268	268	268	268	268	268	268	268
20%	262	320	256	256	384	384	384	384	384	384	384
30%	260	256	256	256	264	264	264	264	264	264	264
40%	4	64	Y	Y	Y	Y	Y	256	256	256	256
50%	4	64	Y	Y	Y	Y	Y	Y	256	256	256
60%	4	64	Y	64	80	80	80	336	336	336	336
70%	4	64	64	80	80	80	112	368	368	368	368
80%	4	64	64	80	80	80	112	368	368	368	368
90%	4	64	64	80	80	80	112	368	368	368	368
100%	4	64	64	80	80	80	112	368	368	368	368

Legend
Y = Satisfies the psychological data
4 = No analogies
64 = SF ≤ AN
80 = LS ≤ AN , SF ≤ AN
112 = LS ≤ AN , SF ≤ AN , LS ≤ FOR
256 = LS ≤ DT 260 = No analogies, LS ≤ DT, LS ≤ DT
262 = No surface matches, no analogies, LS ≤ DT, LS ≤ DT
264 = LS ≤ SF, LS ≤ DT, LS ≤ DT
268 = No analogies, LS ≤ SF, LS ≤ DT, LS ≤ DT
270 = No surface matches, no analogies, LS ≤ SF, LS ≤ DT, LS ≤ DT 3
36 = LS ≤ AN, SF ≤ AN, LS ≤ DT, LS ≤ DT
368 = LS ≤ AN , SF ≤ AN , LS ≤ FOR , LS ≤ DT
384 = AN ≤ FOR, LS ≤ DT

These results are obtained using content vectors which only included relations and logical connectives, leaving out attributes. As before, rows are the width of the MAC selector, Columns correspond to the width of the FAC selector. The codes describe whether that pair of selector widths allows MAC/FAC to account for the human data, and if not, what criteria were violated.

Table 10 Manipulation: Relation-only vectors

These experiments provide evidence that neither attribute information nor relational structure, by themselves, provide the right kind of information to allow the MAC/FAC model to plausibly satisfy the psychological data. Although such generalizations must be viewed with caution, the analysis of why these alternatives fail may be applied to any retrieval model, not just MAC/FAC. Using attribute information alone does not allow a retrieval system to satisfy the rare insights criterion, since the relational information is not used as a cue in retrieval. Using relational information alone tends to violate the scalability criterion, since large fractions of memory must be searched when the discrimination provided by the relational vocabulary is inadequate.

6. Comparing MAC/FAC and ARCS on ARCS datasets

As mentioned earlier, the model of similarity-based retrieval that is closest to MAC/FAC is ARCS (Thagard, et. al. 1990). The ARCS algorithm is shown in Figure 8. ARCS uses a localist connectionist network to apply semantic, structural, and pragmatic constraints to selecting items from memory. Most of the work in ARCS is carried out by the constraint satisfaction network, which provides an elegant mechanism for integrating the disparate constraints that Thagard et al posulate as important to retrieval. The use of competition in retrieval is designed to reduce the number of candidates retrieved. Using pragmatic information provides a means for the system's goals to affect the retrieval process.

Given a pool of memory items I1..In and a probe P:

1. For each item Ii, include it in a matching network if there are any predicates in Ii that are semantically similar to a predicate in P. The matching network implements semantic and structural constraints.
2. Create inhibitory links between units representing competing retrieval hypotheses, to ensure competitive retrieval.
3. Install pragmatic constraints by creating excitatory links between a special pragmatic node and every predicate marked by the user as important.
4. Run the network until it settles.

Figure 8: The ARCS algorithm

After the network settles, an ordering can be placed on nodes representing retrieval hypotheses based on their activation. Unfortunately, no formal criterion was ever specified by which a subset of these retrieval hypotheses is selected to be considered as what is retrieved by ARCS. Consequently, in the experiments below we mainly focus on the subset of retrieval nodes mentioned by Thagard et al in their paper.

6.1 Theoretical tradeoffs

Both models have their appeals and drawbacks. Here we briefly examine several of each.

- *Pragmatic effects:* In MAC/FAC it is assumed that pragmatics and context affect retrieval according to what is encoded in the probe. That is, we assume that plans and goals are important enough to be explicitly represented, and hence will affect retrieval. In ARCS additional influence can be placed on particular subsets of such information by the user marking it as important. The tradeoff between these alternatives will best be explored by embedding them in larger, task-oriented simulations, so we do not consider effects of pragmatics further in this paper.
- *Utility of results:* Because MAC/FAC uses SME in the FAC stage, the result of retrieval can include novel candidate inferences. Since the purpose of retrieval is to find new knowledge to apply to the probe, this is a substantial advantage. ARCS could close this gap somewhat by using ACME (Holyoak & Thagard, 1989) as a post processor.
- *Initial filtering:* MAC/FAC's content vectors represent the overall pattern of predicates occuring in a structured description, so that the dot product cheaply estimates overlap. ARCS' committment to creating a network if there is any predicate overlap places more of the retrieval burden on the expensive process of setting up networks. The inclusive rather than exclusive nature of ARCS' initial stage leads to the paradoxical fact that a system in which pragmatic constraints are central must ignore CAUSE, IF, and other inferentially important predicates to be tractable.

- *Modeling inter-item effects:* Wharton, Holyoak, Downing, Lange, Wickens, and Melz (1994) have shown that ARCS can model effects of competition between memory items in heightening the relative effect of structural similarity to the probe.

Perhaps the most important issue is the notion of *semantic similarity*. A key issue in analogical processing is what criterion should be used to decide if two elements can be placed into correspondence. The FAC stage of MAC/FAC follows the standard structure-mapping position that *analogy is concerned with discovering identical relational systems*. Thus, other elements can be matched flexibly in service of relational matching: any two entities can be placed in correspondence and functions can be matched non-identically if doing so enables a larger structure to match. But relations have only three choices: they can match identically, as in (a); they can fail to match, as in (b); if the surrounding structural match warrants it, they can be re-represented in such a way that part of their representation now matches identically, as in the shift from (c) to (d).

(a) HEAVIER [camel, cow] --- HEAVIER [giraffe, donkey]
(b) HEAVIER [camel, cow] --- BITE [dromedary, calf]
(c) HEAVIER [camel, cow] --- TALLER [giraffe, donkey]
(d) GREATER [WEIGHT(camel), WEIGHT(cow)] --- GREATER [HEIGHT(camel), HEIGHT(cow)].

ACME and ARCS also share the intuition that analogy is a kind of compromise between similarity of larger structures and similarity of individual elements -- *semantic similarity*, in Holyoak and Thagard's (1989) terms. But the local similarity metric is different. These systems use graded similarity at all levels and for all kinds of predicates; relations have no special status. Thus ARCS and ACME might find pair (b) above more similar than pair (a), because of the object similarity. This would not be true for SME and MAC/FAC.

In ACME, semantic similarity was operationalized using similarity tables. For any potential matching term, a similarity table was used to assign a similarity rating, which was then combined with other evidence to decide whether the two predicates could match. Thus, in the examples above, both pair (b) and pair (c) stand a good chance of being matched, depending on the stored similarities between TALLER, HEAVIER, and BITE, camel, dromedary and giraffe, and so on.

In ARCS, an augmented subset of WordNet (Miller, Fellbaum, Kegl, & Miller, 1988) was used to make semantic similarity decisions. WordNet is a psycholinguistic database describing relationships between words. Two predicates in ARCS are considered semantically similar if their corresponding lexical concepts in WordNet are connected via links that denote particular relationships. The use of WordNet as a database for simple lexical inferences is an appealing idea. The lexical connections found in this way should have well-founded motivations. Nevertheless, it important to remember that WordNet was intended as a lexicon, not a language of thought. Using the lexical concepts of WordNet as a predicate vocabulary requires assuming that there exist conceptual representations that correspond to these lexical concepts. That does not seem an implausible assumption. However, assuming that relationships between words, such as *synonym* or *antonym* are used in the cognitive processing of internal representations seems implausible.

We prefer our tiered identicality account, which uses inexpensive inference techniques to suggest ways to re-represent non-identical relations into a canonical representation language. Such canonicalization has many advantages for complex, rich knowledge systems, where meaning arises from the axioms that predicates participate in. When mismatches occur in a context where it is desirable to make the match, we assume that people make use of techniques of re-representation. An example of an inexpensive inference technique to suggest rerepresentation is Falkenhainer's (1987, 1990a) *minimal ascension* method, which looks for common superordinates when context suggested that two predicates should match. The use of

pure identicality augmented by minimal ascension allowed Falkenhainer's PHINEAS system to model the discovery of a variety of physical theories by analogy. We believe that WordNet could be used in a similar fashion, since it has superordinate information.

Holyoak & Thagard have argued that broader (i.e., weaker) notions of semantic similarity are crucial in retrieval, for otherwise we would suffer from too many missed retrievals. Although this at first sounds reasonable, there is a counterargument based on memory size. Human memories are far larger than any cognitive simulation yet constructed. In such a case, the problem of false positives (i.e., too many irrelevant retrievals) becomes critical. False negatives are of course a problem, but they can be overcome to some extent by reformulating and re-representing the probe, treating memory access as an iterative process interleaved with other forms of reasoning (as in Lange & Wharton's (1992a,b, 1993) REMIND model). Thus it could be argued that strong semantic similarity constraints, combined with re-representation, are crucial in retrieval as well as in mapping.

How do these different accounts of semantic similarity fare in predicting patterns of retrieval? In the rest of this section we tackle this question by comparing the performance of MAC/FAC and ARCS on a variety of examples.

6.2 Computational Experiments comparing MAC/FAC and ARCS

6.2.1 Methods

Each experiment below has a similar structure. First each simulation is given a memory, consisting of one or more databases drawn from the ARCS representations.[14] Then retrieval is tested with probes drawn from a small predefined set of stories, replicating Thagard et.al.'s experiments. The memory a simulation operates over consists of one or more databases. In some cases the memory is augmented by a particular story: e.g., when probing with variant Hawk stories, the Thagard et.al. encoding of the "Karla the Hawk" story is added to memory. (This is done to see if the retrieval system is able to find the base story amidst the distractors, given variations on the story as probes.)

For brevity we specify the probe set and memory contents symbolically, using "/" to distinguish probe set from memory and "+" to indicate set union. Thus HAWK/(PLAYS+Karla Base) indicates an experiment where the database of plays was probed with the Hawk stories. A description of the datasets used and these conventions is summarized in Figure 9.

[14] To date we have been unsuccessful in getting ARCS to run on many of the representations we used in Sections 4 and 5. In some cases ARCS' network does not settle after even 1,000 iterations, and run times of up to twelve hours have been required.

Databases:
FABLES = 100 encodings of Aesop's fables, encoded by Thagard et.al.
PLAYS = 25 encodings of Shakespeare's plays, encoded by Thagard et.al.
Story sets used as probes and memory items:
HAWK = Thagard et.al.'s encoding of the "Karla the Hawk" story set, i.e., original story, analog, appearance match, false analogy, and literal similarity versions. Databases using these probes have the original story added to memory, except when the original story itself is used as a probe.
SG = Thagard et.al.'s encoding of the Sour Grapes fable plus variations, i.e., original story, analog, appearance, and literal similarity versions. Databases using these probes have the original story added to memory, except when the original story itself is used as a probe.
H&WSS = Thagard et.al's encoding of Hamlet and West Side Story. When Hamlet is used as a probe it is removed from memory. West Side Story is never placed in memory.
Convention: For convenience, we refer to an experimental setup by the probe stories followed by the database used, e.g., SG/(FABLES+PLAYS) means that the Sour Grapes fables were used as probes with a memory consisting of both plays and fables. When a story is used as a probe, it is removed from memory first.

Figure 9: Databases and experimental stories used in the experiments

Both MAC/FAC and ARCS take propositional representations as inputs, but their representation conventions are quite different. The most crucial difference is that structure-mapping treats attributes, relations, and functions differently, whereas ARCS does not distinguish them. We used the following rules in translation: (1) One-place predicates were classified as attributes, (2) multi-argument predicates were classified as relations, and (3) since the arguments to CAUSE could be either events or modal propositions, we treated predicates used as arguments to a CAUSE statement either as modal relations (e.g., `BECOMING-TRUE`) or functions (e.g., `MARRIED`, `KILLED`). Since functions can be substituted under structure-mapping's identicality critierion, we ran these experiments on representations translated both with and without rule (3) above, i.e., with and without functions. With one exception, noted below, the results were essentially identical with either translation scheme.

All runtimes are measured according to the Lucid Common Lisp internal clock. A single computer[15] was used for both simulations, so that run times would be comperable.

Replication of computational experiments is still something of a novelty, and standards for ensuring that reported simulation results are repeatable have not yet been established in cognitive science. Nevertheless, we have taken many precautions to ensure that we have run ARCS correctly. Where numerical information was available, for instance, we matched numerical results reported by them to several decimal places. One concern was what should count as a retrieval in ARCS. Neither the original ARCS paper nor the code defines a criterion for distinguishing when an item is actually retrieved (indeed, stories with negative activations were sometimes considered retrievals). In reporting ARCS results we cut off the list of retrieved results where Thagard et al. did. In some cases (e.g., fables) this represented a sharp boundary, in other cases (e.g., plays) it did not.

6.3 Experiment 1: Sour Grapes Comparison

In the first study the memory set consists of the fables, including the Sour Grapes fable, and the probes are variants of Sour Grapes. For ARCS results, the numbers in parentheses represent the level of activation it

[15] An IBM RS/6000 Model 530, with 128MB of RAM using Lucid Common Lisp 4.01.

computed for that retrieval. For MAC/FAC results, the numbers in parentheses represent the scores computed by the MAC or FAC stages, as marked.

Table 11 shows the results. The results for ARCS match those reported for the simulation by Thagard et.al. The MAC/FAC results are quite similar. Thus both systems successfully retrieve Sour Grapes from a database of fables when given variations of it. However, MAC/FAC is substantially faster. The runtime difference is fairly typical; MAC/FAC tends to be two orders of magnitude faster than ARCS when tested with identical data on the same computer.

ARCS results

Probe	Results	Seconds
Sour Grapes appearance	Sour Grapes (0.28)	120
Sour Grapes, analog	Sour Grapes (0.21)	81
Sour Grapes, literal similarity	Sour Grapes (0.25)	123

MAC/FAC Results

Probe	Results	Seconds
Sour Grapes appearance	FAC: Sour Grapes (0.53) MAC: Sour Grapes (0.56)	0.3
Sour Grapes analog	FAC: Sour Grapes (2.03) MAC: Sour Grapes (0.62)	0.2
Sour Grapes literal similarity	FAC: Sour Grapes (2.03) MAC: Sour Grapes (0.62)	0.2

For ARCS results, the numbers in parentheses represent the level of activation it computed for that retrieval. For MAC/FAC results, the numbers in parentheses represent the scores computed by the MAC or FAC stages, as marked.

Table 11: Results for SG/Fables experiment

6.4 Experiment 2: Effects of additional memory items on retrieval (Sour Grapes)

To check the stability of results under changes in memory contents, we reran Experiment 1, adding the database of 25 Shakespeare plays encoded by Thagard et al. to the fables database. We then tested the simulations to see if they would retrieve Sour Grapes from the database of 125 fables and plays when probed with variations of Sour Grapes. The results are show in Table 12. MAC/FAC's results remain unchanged, except for a small increase in processing time. ARCS, on the other hand, is distracted by the plays in one of the probe conditions. Increasing the memory by 25% has led to different results with ARCS. The results also hint at a possible size bias in ARCS: It appears to prefer larger descriptions in retrieval, at the cost of correct matches.

ARCS Results

Probe	Results	Seconds
Sour Grapes appearance	Sour Grapes (0.28)	327
Sour Grapes, analog	The Taming of the Shrew (0.22), Merry Wives (0.18), [11 stories], Sour Grapes (-0.19)	251
Sour Grapes, literal similarity	Sour Grapes (0.25)	373

MAC/FAC Results

Probe	Results	Seconds
Sour Grapes appearance	FAC: Sour Grapes (0.53) MAC: Sour Grapes (0.56)	0.4
Sour Grapes analog	FAC: Sour Grapes (2.03) MAC: Sour Grapes (0.62)	0.3
Sour Grapes, literal similarity	FAC: Sour Grapes (2.03) MAC: Sour Grapes (0.62)	0.3

Table 12:Results of SG probes, database = Fables + Plays

6.5 Experiment 3: Larger Probe sizes

While the results for MAC/FAC in Experiment 2 are satisfactory, ARCS' seemingly poor performance requires further investigation. Does the relative size of the probe matter in the memory swamping effect? To find this out, we again ran both simulations, first with the plays database as memory, then with the 25 plays and 100 fables as memory, this time using as probes the Hamlet and West Side Story encodings as probes, as represented by Thagard et.al. Given Hamlet as a probe, the question is whether the systems can retrieve a tragedy, or at least another Shakespeare play. Given West Side Story as a probe, the challenge is more specific: To retrieve Romeo & Juliet, the analogous play.

Table 13 shows the results for plays only in memory, and Table 14 shows the results with both plays and fables in memory. The good news for ARCS is that the fables have only minimally intruded on the activation for the top ranked retreived plays. A Midsummer Night's dream is ARCS' top-ranked retrieval for West Side Story, but it did also, as stated by Thagard et al, retrieve Romeo & Juliet.

ARCS results

Probe	Results	Seconds
Hamlet	Romeo & Juliet (0.54), King Lear (0.53), Othello (0.46), Cymbeline (0.42), Macbeth (0.41), Julius Caesar (0.38)	1843
West Side Story	Midsummer Night's Dream (0.58), Romeo & Juliet (0.57)	2539

MAC/FAC results

Probe	Results	Seconds
Hamlet	FAC: Romeo & Juliet (6.79) MAC: Othello (0.86), Macbeth (0.85), Romeo & Juliet (0.83), Julius Caeser (0.81)	22
West Side Story	FAC: Romeo & Juliet (16.51) MAC: Romeo & Juliet (0.88)	13

Table 13:Results for Hamlet, West Side Story as probes, Plays database.

MAC/FAC, on the other hand, only retrieves Romeo & Juliet with either probe. For West Side Story this is indeed the expected result (and we believe more intuitive than ARCS' result), but what is happening with Hamlet? Examining the structural evaluation scores (e.g., the FAC scores) reveals that FAC considers the match between West Side Story and Romeo & Juliet to be excellent (16.51), which makes sense because the encodings of West Side Story and Romeo & Juliet have almost isomorphic structure. When Hamlet is the probe , FAC is relatively indifferent, the FAC scores were: Romeo & Juliet (6.79), Julius Caesar (5.49), Macbeth (3.72), Othello (2.67). The dropoff from Romeo & Juliet is 20%, which is below than MAC/FAC's default cutoff of 10%.

ARCS Results

Probe	Results	Seconds
Hamlet	Romeo & Juliet (0.531), King Lear (0.528), Othello (0.45), Cymbeline (0.41), Macbeth (0.40), Julius Caesar (0.37)	4112
West Side Story	Midsummer Night's Dream (0.58), Romeo & Juliet (0.57)	5133

MAC/FAC Results

Probe	Results	Seconds
Hamlet	FAC: Romeo & Juliet (6.79) MAC: Othello (0.86), Macbeth (0.85), Romeo & Juliet (0.83), Julius Caesar (0.81), Fable52 (0.80)	26
West Side Story	FAC: Romeo & Juliet (16.51) MAC: Romeo & Juliet (0.88)	8

Table 14:Results for Hamlet, West Side Story as probes, Plays + Fables database.

6.6 Experiment 4: Hawk stories

The goal of encoding the Hawk stories was to replicate the results of Karla the Hawk studies described in Section 2.1.1. Thagard et.al. encoded one story set, and using the relative activation levels of the stories computed by ARCS as relative retrieval probabilities for human subjects. As Section 4.4 pointed out, ARCS' order of retrieval was: literal similarity, first-order overlap, appearance, analogy, which is not a close match to the observed human ordering of literal similarity, appearance, analogy, first-order overlap. By contrast, MAC/FAC matched the human ordinal results in our simulation of this experiment.

However, our purpose in this experiment is to pursue the question of stability of results under different distractors. We ask two questions: (1) does MAC/FAC, using Thagard et.al's encodings, perform appropriately, and (2) does changing the database used as ARCS' memory change its predicted outcomes? Both simulations were run with the Hawk stories as probes, and with the fables (plus Karla story) as memory and with both fables and plays (plus the Karla story) as memory. The results are shown in Table 15 and Table 16 respectively.

ARCS Results

Probe	Results	Seconds
Karl, literal similarity	"Karla" base (0.67)	315
Karla appearance	Fable55 (0.4), [7 fables], "Karla" base (-0.17)	176
Karla, analogy	Fable23 (0.33), [7 fables], "Karla" base (-0.27)	127
Karla, first-order overlap	Fable23 (0.0907), Fable55 (0.0903), [13 fables], "Karla" base (-0.11)	17

MAC/FAC Results.

Probe	Results	Seconds
Karla, Literal Similarity	FAC: "Karla" (16.07) MAC: "Karla" (0.81), Fable71 (0.74)	6
Karla, apperance	FAC: "Karla" (7.92) MAC: "Karla" (0.71), Fable52 (0.71), Fable71(0.66), Fable27(0.65), Fable5(0.64)	7
Karla, analog	FAC: "Karla" (8.57) MAC: "Karla"(0.81), Fable52 (0.77), Fable5 (0.77), Fable71(0.76), Fable45(0.75), Fable59(0.75), Fable27(0.75)	14
Karla, First-order overlap	FAC: "Karla" (5.33), Fable5 (5.33) MAC: "Karla" (0.73), Fable71(0.71), Fable52(0.71), Fable5(0.71), Fable45(0.69), Fable59(0.68),Fable27(0.68)	7

Table 15: Results for HAWK probes, database = Fables + "Karla" base story

No matter which database is used, MAC/FAC always retrieves the Karla story, irrespective of which variant story is used as a probe. The MAC scores explain why: In each case the Karla story is at the top of the ranking, indicating that the pattern of identical predicates overlapping is greater for Karla and variant than for any other story. The fact that the Karla base story is retrieved for the literal similarity and appearance variants is expected. Its retrieval when the analogy is used as a probe is also reasonable (although if ARCS always retrieved analogs successfully it would be an implausible model). Retrieving the

base story when the first-order overlap story is used as a probe is not so reasonable. We believe this occurs because the Thagard et.al. representations are rather sparse and include almost no surface information, and thus are less natural than might be desired (c.f. the *specificity conjecture* of Forbus & Gentner, 1989).

ARCS Results

Probe	Results	Seconds
"Karla", literal similarity	"Karla" base (0.67)	614
"Karla", appearance	Fable55 (0.40),[16 stories], "Karla" base (-0.018)	408
"Karla", analogy	Pericles (0.60), [17 stories], "Karla" base (-0.32)	244
"Karla", false analogy	Pericles (0.58), [22 stories], "Karla" base (-0.38)	45

MAC/FAC Results

Probe	Results	Seconds
Karla, Literal Similarity	FAC: "Karla"(16.07) MAC: "Karla"(0.81), Fable71 (0.74)	7
Karla, apperance	FAC: "Karla" (7.92), MAC: "Karla" (0.71), Fable52(0.71), Julius Caesar (0.69), Othello (0.68), Macbeth (0.67), Fable71(0.66), Two Gentlemen of Verona (0.65), Fable27(0.65), Hamlet (0.65), Fable5(0.64)	21
Karla, analog	FAC:"Karla"(8.57) MAC: "Karla" (0.81), Julius Caesar (0.78), Two Gentlemen of Verona (0.78), Fable52(0.77), Fable5(0.77), Macbeth (0.76),As You Like It(0.76), Fable71(0.76), Fable45(0.75), Fable59(0.75), Fable27(0.75), Othello(0.75)	37
Karla, First-order overlap	FAC: "Karla"(5.33), Fable5(5.33) MAC: "Karla"(0.73), Juilius Caesar(0.72), Two Gentlemen of Verona (0.72), Fable71(0.71), Fable52(0.71), Fable5 (0.71), Macbeth(0.70), As You Like It (0.70), Othello (0.69), Fable45 (0.69), Hamlet(0.68)	23

Table 16: Results for HAWK probes, with database = Fables + Plays + "Karla" base story

Interestingly, this experiment marks the only place where the decision to use functions in encoding made any real difference in the results. If no functions were used in translating the ARCS representations, the MAC results remained the same (because content vectors are based strictly on identical predicates), but the Karla base story would be knocked out of the FAC output by other stories that had more overlapping structure, since the causal structure in the Karla story could not be consistently mapped due to non-identical relations. The fact that this problem only shows up with this one probe set, out of all the representations made by Thagard et.al., suggests that this is not a serious problem.

As was suggested by experiments 1 and 2, the ARCS results vary considerably with different distractor sets. This means that the use of relative activations to estimate relative frequencies is not a stable measure. Specifically, the relative ordering of first-order overlap and analogy reverses when the database of fables is augmented with the plays. The position of the Karla story in the activation rankings is also alarming. The appearance story, which should retrieve the base almost as often as the literal similarity story, has dropped from ninth in the ranking to 18th. Depending on where the retrieval cutoff is placed, the conclusion might be that ARCS fails to retrieve the Karla story given the very close surface match.

6.7 Experiment 5: ARCS using identicality

The results so far indicate that MAC/FAC is far more immune to false positives than ARCS. What is responsible for this difference? Is it MAC/FAC's use of a separate stage that performs structural filtering? The use of content vectors versus parallel constraint satisfaction to generate an initial set of retrieval candidates? MAC/FAC's identicality constraint versus ARCS' weaker semantic similarity constraint? A complete answer to this question will require much more empirical and theoretical work, but we can gain some insight by a simple experiment. We ran ARCS again, but without the WordNet-inspired similarity network. Under such conditions, ARCS only creates local matches between identical predicates, and the initial candidate set is much smaller, because the semantic similarity constraint has been greatly tightened.

The results of this experiment are shown in Tables 17-19. Table 17 shows that the results on Sour grapes have improved substantially; ARCS is no longer tempted by plays. Table 18 shows that, while a Midsummer Night's Dream is high on ARCS' list, it no longer prefers it to Romeo & Juliet when West Side Story is used as a probe. The HAWK results show the least improvement; the estimated retrieval order again does not match that of human subjects, and there are still many fables and plays ahead of what should be very close matches to the Karla base story.

ARCS w/identicality, SG/FABLES

Probe	Results	Seconds
Sour Grapes literal similarity	Sour Grapes(0.18)	1.3
Sour Grapes appearance	Sour Grapes (0.28)	23
Sour Grapes analogy	Sour Grapes (0.18)	1.1

ARCS w/identicality, SG/(FABLES+PLAYS)

Probe	Results	Seconds
Sour Grapes literal similarity	Sour Grapes (0.19)	4
Sour Grapes appearance	Sour Grapes (0.28)	34
Sour Grapes analogy	Sour Grapes (0.19)	4

Table 17: ARCS w/identicality on Sour Grapes with Fables and Fables + Plays

ARCS w/identicality, database = plays

Probe	Result	Seconds
Hamlet	King Lear (0.56), Romeo & Juliet (0.52), Othello (0.47), Cymbeline (0.41), Macbeth (0.40), Julius Caesar (0.38)	489
West Side Story	Romeo & Juliet (0.59), Midsummer Night's Dream (0.52)	1671

ARCS w/identicality, database = plays+Fables

Probe	Result	Seconds
Hamlet	King Lear (0.55), Romeo & Juliet (0.51), Othello (0.46), Cymbeline (0.49), Macbeth (0.39), Julius Caesar (0.37)	1108
West Side Story	Romeo & Juliet (0.59), Midsummer Night's Dream (0.52)	3014

Table 18: ARCS w/identicality, probed with Plays

ARCS w/identicality, HAWK/FABLES

Probe	Results	Seconds
Karla, Literal Similarity	Fable23(0.261),Fable55(0.258),Karla story (-0.1)	73
Karla, Appearance	Fable55(0.4),[8 fables],Karla story (-0.23)	114
Karla, True Analogy	Fable23(0.26),Fable55(0.26),[5 fables], Karla story (-0.23)	12
Karla, First-Order overlap	Fable23(0.087),Fable55(0.087)	5

ARCS w/identicality, HAWK/(FABLES+PLAYS)

Probe	Results	Seconds
Karla, Literal Similarity	Fable23(0.26),Fable55(0.26),Karla story(-0.014)	74
Karla, Appearance	Fable55(0.25),Hamlet(0.17),Fable23(0.067),[17 plays & fables],Karla story(-0.22)	154
Karla, True Analogy	Pericles(0.55),[3 plays] Fable23(-0.13),Fable55(-0.13), [8 plays & fables],Karla story (-0.30)	29
Karla, First-Order overlap	Pericles(0.59),[6 fables & plays],Fable23(-0.25), Fable55(-0.25)	18

Table 19: ARCS w/identicality, HAWK probes

6.8 Conclusions from Computational Comparison Experiments

The results of cognitive simulation experiments must always be interpreted with care. In this case, we believe our experiments provide evidence that MAC/FAC, using structure-mapping's identicality constraint, better models retrieval than ARCS, which uses Thagard et al.'s notion of semantic similarity. In retrieval, the special demands of large memories argue for simpler algorithms, simply because the cost of false positives is much higher. If retrieval were a one-shot, all-or-nothing operation, the cost of false negatives would be higher. But that is not the case. In normal situations, retrieval is an iterative process, interleaved with the construction of the representations being used. Thus the cost of false negatives is reduced by the chance that reformulation of the probe, due to rerepresentation and inference, will subsequently catch a relevant memory that slipped by once.

While both are designed to allow parallel implementations, the two orders of magnitude speed difference clearly suggests that MAC/FAC is the more practical choice for cognitive simulation experiments.

Finally, we note that while ARCS' use of a localist connectionist network to implement constraint satisfaction is in many ways intuitively appealing, it is by no means clear that such implementations are neurally plausible. On the other hand, we believe the evidence suggests that MAC/FAC captures similarity-based retrieval phenomena better than ARCS does.

7. Discussion

Similarity is central in transfer. To understand its role requires making fine distinctions both about similarity and about transfer. The psychological evidence indicates that the accessibility of matches from memory is strongly influenced by surface commonalities and weakly influenced by structural commonalities, while the rated inferential soundness of comparisons is strongly influenced by structural commonalities and little, if at all, influenced by surface commonalities. An account of similarity in transfer must deal with this disassociation between retrieval and structural alignment: between the matches people *get* from memory and the matches they *want*.

The MAC/FAC model of similarity-based retrieval captures both the fact that humans successfully store and retrieve intricate relational structures and the fact that access to these stored structures is heavily (though not entirely) surface-driven. The first stage is attentive to content and blind to structure and the second stage is attentive to both content and structure. The MAC stage uses content vectors, a novel summary of structured representations, to provide an inexpensive "wide net" search of memory, whose results are pruned by the more expensive literal similarity matcher of the FAC stage to arrive at useful, structurally sound matches.

The simulation results presented here demonstrate that MAC/FAC can simulate the patterns of access exhibited by human subjects. It displays the appropriate preponderance of literal similarity and surface matches, and it occasionally retrieves purely relational matches (Section 4). Our sensitivity studies suggest that these results are a consequence of our theory, and are not hostage to non-theoretically motivated parameters or algorithmic choices (Section 5). Our computational experiments comparing MAC/FAC and ARCS (Section 6) suggests that MAC/FAC both accounts for the psychological results more accurately and more robustly than ARCS. In addition to the experiments reported here, we have tested MAC/FAC on a variety of other data sets, including relational metaphors (30 descriptions, average of 12 propositions each) and attribute-rich descriptions of physical situations as might be found in commonsense reasoning (12 descriptions, averaging 42 propositions each). We have also tried various combinations of these databases with the Karla the Hawk data set (45 descriptions, averaging 67 propositions each). In all cases

to date MAC/FAC's performance has been satisfactory and consistent with the overall pattern of findings regarding human retrieval.

However, there are still many open issues regarding access and MAC/FAC. The rest of this section examines the most important of these, compares MAC/FAC to other directly relevant models, and explores some broader implications.

7.1 Limitations and open questions

7.1.1 Retrieval failure

Sometimes a probe reminds us of nothing. Currently the only way this can happen in the MAC/FAC model is for FAC to reject every candidate provided by MAC. This can happen if no structurally sound match hypotheses can be generated between the probe and the descriptions output by MAC. (Without any local correspondences there can be no interpretation of the comparison.) This can happen, albeit rarely. A variant of MAC/FAC with thresholds on the output of either or the both MAC and FAC stages -- so that the system would return nothing if the best match were below criterion -- would show more non-remindings.

7.1.2 Focused remindings and penetrability

Many AI retrieval programs and cognitive simulations elevate the reasoner's current goals to a central role in their theoretical accounts (e.g., Hammond, 1986, 1989; Keane, 1988ab; Kolodner, 1984, 1989; Riesbeck & Schank, 1989; Thagard, Holyoak, Nelson, & Gochfeld, 1990). Although we agree with the claim that goal structures are important, MAC/FAC does not give goals a separate status in retrieval. Rather, we assume that the person's current goals are represented as part of the higher-order structure of the probe. The assumption is that goals are embedded in a relational structure linking them to the rest of the situation; they play a role in retrieval, but the rest of the situational factors must participate as well. When one is hungry, for instance, presumably the ways of getting food that come to mind are different if one is standing in a restaurant, a supermarket, or in the middle of a forest. The inclusion of current goals as part of the representation of the probe is consistent with the finding of Read and Cesa (1991) that asking subjects for explanations of current scenarios leads to a relatively high rate of analogical reminding. However, we see no reason to elevate goals above other kinds of higher-order structure. By treating goals as just one of many kinds of higher-order structures, we escape making the erroneous prediction of many case-based reasoning systems: that retrieval requires common goals. People can retrieve information that was originally stored under different goal structures. (See Goldstein, Kedar, & Bareiss, 1993, for a discussion of this point.)

A related question concerns the degree to which the results of each stage are inspectable and tunable. We assume that the results of the FAC stage are inspectable, but that explicit awareness of the results of the MAC stage is lacking. We conjecture that one can get a sense that there are possible matches in the MAC output, and perhaps some impression of how strong the matches are, but not what those items are. The feeling of being reminded without being able to remember the actual item might correspond to having candidates generated by MAC which are all either too weak to pass on or are rejected by the FAC stage.

How much can MAC and FAC be affected by the subject? There is psychological evidence that people cannot directly control the kinds of matches they retrieve. Schumacher and Gentner (in preparation) investigated this by varying the test instructions given to subjects. They gave subjects lists of proverbs to

read, followed by test proverbs that were either structurally similar or surface-similar to proverbs studied previously. Subjects who were told to write any prior proverbs that they were reminded of while reading the test proverbs recalled about twice as many surface matches as analogies. Another group of subjects was told to write only structural remindings and to strive for as many of these as possible. Although these subjects indeed wrote many fewer surface matches than the first group, they recalled only the same low number of analogies. The goal to seek relational matches apparently led people to filter non-relational matches, but not to find more relational matches. This suggests that the FAC matcher may be tunable (in that subjects were able to filter out the surface matches) but not the MAC matcher (in that subjects were not able to produce more analogies on demand).

The idea that FAC, though not MAC, is tunable is consistent with evidence that people can be selective in similarity matching once both members of a pair are present. For example, in a triads task, matching XX to OO or XO, subjects can readily choose either only relational matches (XX -- OO) or only surface matches (XX -- XO) (Gentner & Markman, 1994a,b; Goldstone, Medin & Gentner, 1991; Medin, Goldstone & Gentner, 1993). This could be modeled by assuming either that the FAC matcher is tuned according to such task constraints, or that there is a separate tunable similarity processor for comparing items within working memory.

7.1.3 Size of content vectors

One potential problem with scaling up MAC/FAC is the potential growth in the size of content vectors. Our current descriptions use a vocabulary of only a few hundred distinct predicates. We implement content vectors via sparse encoding techniques, analogous to those used in computational matrix algebra, for efficiency. However, a psychologically plausible representation vocabulary may have hundreds of thousands of predicates. It is not obvious that our sparse encoding techniques will suffice for vocabularies that large, nor does this implementation address the question of how systems with limited "hardware bandwidth", such as connectionist implementations, could serve as a substrate for this model.

This scale problem is mitigated partly by MAC/FAC's basic architecture with its cheap initial filter. However there are at least two further possible ways to address the potential scale problem in the size of content vectors. The first is *abstraction*. In symbolic knowledge representations, predicates and functions are often arranged in hierarchies. For example, a complex concept such as *bequeath* might be stored as a specialization of the concept of *giving*, which might in turn be a specialization of the concept of *transfer*. Let us view the set of specializations between predicates as a lattice. Any set of predicates that partitions the lattice can be used to formulate a *semantically compressed* content vector as follows: The weight of a component of the compressed content vector is a function of the number of occurrences of that predicate -- and all predicates below it in the partition -- in the description. In effect, predicates below the selected subsets are replaced with more abstract versions. Another possible solution for the scale problem is *factorization*: the predicates could be partitioned into subsets that are tightly interrelated and separate content vectors computed for each subset. This organization presumes that there is some fixed size bound on processing modules, but that several processing modules can be synchronized well enough to accumulate results across them.

7.1.4 Combining similarity effects across items

MAC/FAC is currently a purely exemplar-based memory system. The memory items can be highly situation-specific encodings of perceptual stimuli, abstract mathematical descriptions, causal scenarios, etc, MAC/FAC lacks the capacity to inter-item effects. For example, MAC/FAC does not capture competition among items. Wharton, Holyoak, Downing, Lange, and Wickens (1991, 1992) and Wharton, Holyoak,

Downing, Lange, Wickens, and Melz (1994) have shown an intriguing effect where competition between exemplars heightens the relative effect of structural similarity in retrieval. MAC/FAC also does not average across several items at retrieval (Medin & Schaffer, 1978) or derive a global sense of familiarity by combining the activations of multiple retrievals (Gillund & Shiffrin, 1984; Hintzman, 1986, 1988). An interesting extension of MAC/FAC would be to include this kind of between-item processing upon retrieval.

If such inter-item averaging occurs, it could provide a route to the incremental construction of abstractions and indexing information in memory. We see three plausible ways to do this. First, as above, the descriptions output by the MAC stage could be compared. Second, the access system might incrementally build up something like Minsky's (1981) similarity network, using the history of retrievals to encode difference descriptions to simplify future access. Third, the descriptions output by the FAC stage could be compared: SME could be used to carry out structural abstraction across several descriptions (as in Skorstad, Medin, & Gentner, 1988) to produce a combined description as the FAC output. The first and third models are both forms of "late averaging" accounts, and it would be interesting to compare these techniques with other models that account for prototype effects by combining exemplars at retrieval (Hintzman, 1986, 1988; Medin & Shaffer, 1978).

7.1.5 Iterative access

Keane (1988c, 1991; Keane & Brayshaw, 1988) and Burstein (1983a,b) have proposed incremental mapping processes. We suggest that similarity-based retrieval may also be an iterative process. In particular, in active retrieval (as opposed to spontaneous remindings) we conjecture that MAC/FAC may be used iteratively, each time modifying the probe in response to the previous match (c.f. Falkenhainer 1987, 1990; Gentner 1989). Suppose for example a probe yielded several partial remindings. The system of matches could provide clues as to which aspects of the probe are more or less relevant, and thus should be highlighted or suppressed on the next iteration. MAC should respond to this altered vector by returning more relevant items, and FAC can then select the best of these.

Another advantage of such incremental reminding is that it might help explain how we derive new relational categories. Barsalou's (1982, 1987) *ad hoc* categories, such as "things to take on a picnic" and Glucksberg and Keysar's (1990) metaphorically based categories, such as "jail" as a prototypical confining institution, are examples of the kinds of abstract relational commonalities that might be highlighted during a process of incremental retrieval and mapping.

7.1.6 Embedding in performance-oriented models

MAC/FAC is not itself a complete analogical processing system. For example, both constructing a model from multiple analogs (e.g., Burstein, 1983a,b) and learning a domain theory by analogy (e.g., Falkenhainer 1987, 1988, 1990b) require multiple iterations of accessing, mapping, and evaluating descriptions. Several psychological questions about access cannot be studied without embedding MAC/FAC in a more comprehensive model of analogical processing. First, as discussed above, there is ample evidence that subjects can choose to focus on different kinds of similarity when the items being compared are both already in working memory. Embedding MAC/FAC in a larger system should help make clear whether this penetrability should be modeled as applying to the FAC system or to a separate similarity engine. (Order effects in analogical problem solving (Keane, in press) suggest the latter.)

A second issue that requires a larger, performance-oriented model to explore via simulation is when and how pragmatic constraints should be incorporated (c.f. Holyoak & Thagard, 1989; Thagard et al, 1990).

Since we assume that goals, plans, and similar control knowledge is explicitly represented in working memory, the MAC stage will include such predicates in the content vector for the probe, and hence will be influenced by pragmatic concerns. There are two ways to model the effects of pragmatics on the FAC stage. The first is to use the SME *pragmatic marking algorithm* (Forbus and Oblinger, 1990) as a relevance filter. The second is to use incremental mapping, as in Keane & Brayshaw's (1988) Incremental Analogy Machine (IAM). This technique permits the selection and grouping of sets of correspondences to be influenced by the task at hand (Forbus, Ferguson, & Gentner, 1994).

A recent simulation by Lange and Wharton (1992, 1993) called REMIND models retrieval in the context of natural language processing, using spreading activation in a connectionist network both to construct a conceptual representation from textual input and to find the most similar story in its episodic memory. REMIND is an intriguing model, and the attempt to integrate multiple cognitive processes into larger models is an important activity. However, it is difficult to compare this model with MAC/FAC and other retrieval models. First, REMIND only models a specific retrieval task, namely retrieval in the service of understanding stories, and thus does not attempt to cover as wide a span of phenomena as MAC/FAC. Second, when REMIND retrieves a story, it does not appear to create correspondences between the understanding of its input and the previous story, nor does it generate novel candidate inferences, and thus does not satisfy the structured mappings criterion for retrieval. Third, REMIND has only been tested on a corpus involving a handful of short (i.e., two sentence) stories. To our knowledge, it has never been tested either on a corpus as large as those used with MAC/FAC and ARCS, nor on a corpus that includes examples as large as those used with MAC/FAC. Even their current small databases stretch the limits of a Connection Machine,[16] which makes it difficult to evaluate their model thoroughly.

7.1.7 Expertise and Relational Access

Despite the gloomy picture painted in the present research and in most of the problem-solving research, there is evidence of considerable relational access (a) for experts in a domain and (b) when initial encoding of the study set is relatively intensive. Novick (1988a,b) studied remindings for mathematics problems using novice and expert mathematics students. She found that experts were more likely than novices to retrieve a structurally similar prior problem, and when they did retrieve a surface-similar problem, they were quicker to reject it than were novices. Faries and Reiser (1988) taught subjects LISP in a series of intensive training sessions and then gave them target problems that were superficially similar to one prior problem and structurally similar to another. Given this intensive training, Faries and Reiser's subjects were able to access structurally similar problems despite the competing superficial similarities.

The second contributor to relational retrieval, almost certainly related to the first, is intensive encoding. Gick and Holyoak (1983) and Catrambone and Holyoak (1987, 1989) found that subjects exhibited increased relational retrieval when they were required to compare two prior analogs, but not when they were simply given two prior analogs. Schumacher & Gentner (1987) found increased relational retrieval of proverbs when subjects wrote out the meaning of each proverb on the study list, as opposed to simply reading it or rating its cleverness. Seifert, McKoon, Abelson, and Ratcliff (1986) investigated priming effects in a sentence verification task between thematically similar (analogical) stories. They obtained priming when subjects first studied a list of themes and then judged the thematic similarity of pairs of stories, but not when subjects simply read the stories.

The increase of relational reminding with expertise and with intensive encoding can be accommodated in the MAC/FAC model. First, we assume that experts have richer and better structured representations of the

[16] Trent Lange, personal communication, IJCAI-93.

relations in the content domain than do novices (Carey, 1985; Chi, 1978; Reed, Ackinclose, & Voss, 1990). This fits with developmental evidence that as children come to notice and encode higher-order relations such as *symmetry* and *monotonicity*, their appreciation of abstract similarity increases (Kotovsky & Gentner, 1990; Gentner & Rattermann, 1991). Second, in particular we speculate that experts may have a more uniform internal relational vocabulary within the domain of expertise than do novices (Clement, Mawby & Giles, 1994; Gentner & Rattermann, 1991; Gentner, Rattermann, Markman, & Kotovsky, in press). The idea is that experts tend to have relatively comprehensive theories in a domain and that this promotes canonical relational encodings within the domain.

To the extent that a given higher-order relational pattern is used to encode a given situation, it will of course be automatically incorporated into MAC/FAC's content vector. This means that any higher-order relational concept that is widely used in a domain will tend to increase the uniformity of the representations in memory. This increased uniformity should increase the mutual accessibility of situations within the domain. Thus as experts come to encode a domain according to a uniform set of principles, the likelihood of appropriate relational remindings increases. That is, under the MAC/FAC model, the differences in retrieval patterns for novices and experts are explained in terms of differences in knowledge, rather than by the construction of explicit indices.

Bassok has made an interesting argument that indirectly supports this claim of greater relational uniformity for experts than for novices (Bassok & Wu, in preparation). Noting prior findings that in forming representations of novel texts people's interpretations of verbs depend on the nouns attached to them (Gentner, 1981; Gentner & France, 1988), Bassok suggests that particular representations of the relational structure may thus be idiosyncratically related to the surface content, and that this is one contributor to the poor relational access. If this is true, and if we are correct in our supposition that experts tend to have a relatively uniform relational vocabulary, then an advantage for experts in relational access would be predicted.

As domain expertise increases, MAC/FAC's activity may come to resemble a multi-goal case-based reasoning model. We can think of its content vectors as indices with the property that they change automatically with any change in the representation of domain exemplars. Thus as domain knowledge -- particularly the higher-order relational vocabulary -- increases, MAC/FAC may come to have sufficiently elaborated representations to permit a fairly high proportion of relational remindings. The case-based reasoning emphasis on retrieving prior examples and generalizations that are inferentially useful may thus be a reasonable approximation to the way experts retrieve knowledge.

7.1.8 Comparison with recent approaches

Retrieval models tend to be of two types: Cognitive simulations that are committed to modeling human performance in detail, and case-based reasoning systems that, while inspired by human cognition, are equally concerned with effective performance on practical problems. Of recent cognitive simulations, the closest is the REMIND model of Lange and Wharton (1992, 1993), discussed in Section 7.1.6. Of the case-based reasoning systems, one of closest to our approach seems to be the CaPER system (Kettler, Hendler, and Anderson, 1992). CaPER is designed to retrieve all sufficiently similar plans from an unindexed case-base, using massive parallelism to do a simple, non-structural match between a query and the contents of memory. It would be very interesting how well the parallel techniques used in CaPER could be applied to MAC/FAC.

7.2 The Decomposition of Similarity

The dissociation between surface similarity and structural similarity across different processes has broader implications for cognition, and is related to several recent discussions. Medin, Goldstone and Gentner (1993) and Gentner (1989) have argued that similarity is pluralistic, in the sense that there are multiple subclasses of similarity and multiple influences on how it is computed. Rips (1989) demonstrated a dissociation between similarity, typicality, and categorization. Murphy and Medin (1985) and Keil (1989) have commented on the limited usefulness of simple similarity and pointed out that physical resemblance does not provide a sufficient basis for determining conceptual groupings. As discussed above, there is a relational shift in development (Gentner & Rattermann, 1991; Halford, 1992). Finally, local object matches appear to be processed faster by adults than structural commonalities. Goldstone and Medin (1994) found that local similarities have their effects on mapping earlier than global relational similarities in a timed mapping task, and Ratcliff and McKoon (1989) found convergent results in a sentence-matching task: subjects could discriminate new from old sentences faster if the new sentences contained all new words (e.g., "Helen attracted Jeff.". vs. "Andrew accosted Mary.") than if the sentences differed only in relational structure (e.g., "Helen attracted Jeff." vs. "Jeff attracted Helen."). In pilot experiments using perceptual stimuli, in which subjects were timed under different kinds of mapping instructions, Markman and Gentner (in press) have found that subjects are faster to choose on the basis of similar objects than on the basis of similar relations, even when the two rules dictate the same response.

These kinds of results render less plausible the notion of a unitary similarity that governs retrieval, evaluation and inference. Instead, they suggest a more complex, pluralistic view of similarity. MAC/FAC provides an architecture which demonstrates how such a pluralistic notion of similarity can be organized to account for psychological data on retrieval.

8. Acknowledgements

This research was supported by the Office of Naval Research (Contract No. N00014-89-J-1272). We thank Ray Bareiss, Mark Burstein, Gregg Collins, Ron Ferguson, Brian Falkenhainer, Rob Goldstone, Art Markman, Doug Medin, and Mary Jo Rattermann for discussions of these issues. We thank Paul Thagard for providing us with ARCS and its associated databases.

9. References

Barnden, J. A. (1994). On the connectionist implementation of analogy and working memory matching. In K. J. Holyoak & J. A. Barnden (Eds.), *Advances in connectionist and neural computation theory, Vol. 2: Analogical connections*. Norwood, NJ: Ablex Publishing Company, 327-374.

Barsalou, L. W. (1982). Context-independent and context-dependent information in concepts. *Memory and Cognition, 10,* 82-93.

Barsalou, L. W. (1987). The instability of graded structure: Implications for the nature of concepts. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization.* Cambridge, England: Cambridge University Press.

Bassok, M. & Wu, L. (in preparation). Judging a book by its cover: Interpretative effects of content on problem solving transfer.

Becker, L., & Jazayeri, K. (1989). A connectionist approach to case-based reasoning. *Proceedings: Case-based Reasoning Workshop* (pp. 213-217). San Mateo, CA: Morgan Kaufmann.

Birnbaum, L., & Collins, G. (1989). Remindings and engineering design themes: A case study in indexing vocabulary. *Proceedings: Case-based Reasoning Workshop* (pp. 47-51). San Mateo, CA: Morgan Kaufmann.

Branting, K. L. (1991). Building Explanations from Rules and Structured Cases. To appear in *IJMMS.*

Brown, A. L. (1989). Analogical learning and transfer: What develops? In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 369-412). Cambridge, MA: Cambridge University Press.

Burstein, M. H. (1983). Concept formation by incremental analogical reasoning and debugging. *Proceedings of the International Machine Learning Workshop* (pp. 19-25). Monticello, IL. Urbana, IL: University of Illinois.

Burstein, M. H. (1983). A model of learning by incremental analogical reasoning and debugging. *Proceedings of the National Conference on Artificial Intelligence* (pp. 45- 48). Washington, D. C. Los Altos, CA: Morgan Kaufmann.

Burstein, M. (1989). Analogy vs. CBR: The purpose of mapping. *Proceedings: Case-based Reasoning Workshop* (pp. 133-136). San Mateo, CA: Morgan Kaufmann.

Carey, S. (1985). *Conceptual change in childhood.* Cambridge, MA: MIT Press.

Catrambone, R., & Holyoak, K. J. (May, 1987). *Do novices have schemas?* Paper presented at the Fifty-Second Annual Meeting of the Midwestern Psychological Association, Chicago, IL.

Catrambone, R., & Holyoak, K. J. (1989). Overcoming contextual limitations on problem-solving transfer. *Journal of Experimental Psychology: Learning, Memory and Cognition, 15.*

Chi, M. T. H. (1978). Knowledge structures and memory development. In R. S. Siegler (Ed.), *Children's thinking: What develops?* (pp. 73-96). Hillsdale, NJ: Erlbaum.

Clement, J. (1986). Methods for evaluating the validity of hypothesized analogies. *Proceedings of the Eighth Annual Conference of the Cognitive Science Society*, (pp 223-234), Amherst, MA. Hillsdale, NJ: Erlbaum.

Clement, C. A., & Gentner, D. (1991). Systematicity as a selection constraint in analogical mapping. *Cognitive Science, 15,* 89-132.

Clement, C. A., & Mawby, R., & Giles, D. E. (1994). The effects of manifest relational similarity on analogy retrieval. *Journal of Memory and Language 33,* (396-420).

Elio, R., & Anderson, J. R. (1981). The effects of category generalizations and instance similarity on schema abstraction. *Journal of Experimental Psychology: Human Learning and Memory, 7,* 397-417.

Elio, R., & Anderson, J. R. (1984). The effects of information order and learning mode on schema abstraction. *Memory & Cognition, 12,* 20-30.

Falkenhainer, B. (1987). An examination of the third stage in the analogy process: Verification-based analogical learning. *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, August, 1987.

Falkenhainer, B. (1988). *Learning from physical analogies: A study of analogy and the explanation process* (Tech. Rep. No. UIUCDCS-R-88-1479). Urbana, IL: University of Iliinois, Department of Computer Science.

Falkenhainer, B. (1990a). Analogical Interpretation in Context. *Proceedings of the Cognitive Science Society.*

Falkenhainer, B. (1990b). A unified approach to explanation and theory formation. In Shrager, J. and Langley, P. (Eds.) *Computational models of scientific discovery and theory formation*, Morgan-Kaufmann, pp 157-196.

Falkenhainer, B., Forbus, K., Gentner, D. (1989). The Structure-Mapping Engine: Algorithm and examples. *Artificial Intelligence, 41,* 1-63.

Faries, J. M., & Reiser, B. J. (1988). Access and use of previous solutions in a problem solving situation. *Proceedings of the Tenth Annual Meeting of the Cognitive Science Society* (pp. 433-439). Montreal; Hillsdale, NJ: Erlbaum.

Forbus, K., & Gentner, D. (1986). Learning physical domains: Towards a theoretical framework. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach* (Vol. 2, pp. 311-348). Los Altos, CA: Morgan Kaufmann.

Forbus, K., & Gentner, D. (1989). Structural evaluation of analogies: What counts? *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society* (pp. 341-348). Ann Arbor, MI. Hillsdale, NJ: Erlbaum.

Forbus, K. D., & Oblinger, D. (1990). Making SME greedy and pragmatic. *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society* (pp. 61-68). Cambridge, MA. Hillsdale, NJ: Erlbaum.

Gelman & J. P. Brynes (Eds.), *Perspectives on Language and Thought: Interrelations in Development* (pp. 225-277). London: Cambridge University Press.

Forbus, K.D., Ferguson, R. W., and Gentner, D. (1994) Incremental structure mapping. *Proceedings of the Annual Conference of the Cogntiive Science Society*, to appear.

Gentner, D. (1981). Some interesting differences between nouns and verbs. *Cognition and Brain Theory, 4* (2), 161-178.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science, 7*, 155-170.

Gentner, D. (1988a). Analogical inference and analogical access. In A. Prieditis (Ed.), *Analogica* (pp. 63-88). Los Altos, CA: Morgan Kaufmann.

Gentner, D. (1988b). Metaphor as structure mapping: The relational shift. *Child Development, 59*, 47-59.

Gentner, D. (1989a). The mechanisms of analogical learning. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 199-241). New York: Cambridge University Press.

Gentner, D. (1989b). Finding the needle: Accessing and reasoning from prior cases. In *Proceedings: Case-Based Reasoning Workshop*, Defense Advanced Research Projects Agency, Information Science and Technology Office (pp. 137-143).

Gentner, D., & Clement, C. (1988). Evidence for relational selectivity in the interpretation of analogy and metaphor. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 22, pp. 307-358).

Gentner, D., & Landers, R. (1985). Analogical reminding: A good match is hard to find. *Proceedings of the International Conference on Cybernetics and Society* (pp. 607-613). Tucson, AZ.

Gentner, D., & France, I. M. (1988). The verb mutability effect: Studies of the combinatorial semantics of nouns and verbs. In S. L. Small, G. W. Cottrell, & M. K. Tanenhaus (Eds.), *Lexical ambiguity resolution: Perspectives from psycholinguistics, neuropsychology, and artificial intelligence* (pp. 343-382). San Mateo, CA: Morgan Kaufmann.

Gentner, D., & Markman, A. B. (1993). Analogy: Watershed or Waterloo: Structural alignment and the development of connectionist models of cognition. In S. J. Hanson, J. D. Cowan & C. L. Giles (Eds.), *Advances in Neural information Processing Systems 5* (pp. 855-862). San Mateo, CA: Morgan Kaufmann.

Gentner, D., & Markman, A. B. (1994). Structural alignment in comparison: No difference without similarity. *Psychological Science, 5*(3), 153-158.

Gentner, D., & Markman, A. B. (in press). Similarity is like analogy. To appear in C. Cacciari (Ed). *Proceedings of the Workshop on Similarity at the University of San Marino*. Milan, Italy: Bompiani.

Gentner, D., & Rattermann, M. J. (1991). Language and the career of similarity. In S. A. Gelman & J. P. Brynes (Eds.), *Perspectives on Language and Thought: Interrelations in Development* (pp. 225-277).

Gentner, D., Rattermann, M. J. & Forbus, K. D. (1993). The roles of similarity in transfer: Separating retrieval from inferential soundness. *Cognitive Psychology, 25,* 524-575

Gentner, D., Rattermann, M. J., Kotovsky, L., & Markman, A. B. (in press). Two forces in the development of relational similarity. In G. Halford and T. Simon (Eds.), *Developing cognitive competence: New approaches to process modeling.* Hillsdale, NJ: Erlbaum.

Gick, M. L., & Holyoak, K. J. (1980). Analogical problem solving. *Cognitive Psychology, 12,* 306-355.

Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology, 15,* 1-38.

Gillund, G., & Shiffrin, R. M. (1984). A retrieval model for both recognition and recall. *Psychology Review, 91,* 1-67.

Glucksberg, S., & Keysar, B. (1990). Understanding metaphorical comparisons: Beyond similarity. *Psychological Review, 97,* 3-18.

Goldstein, E., Kedar, S., and Bareiss, R. (1993). Easing the creation of a multi-purpose case library. *Proceedings of the AAAI-93 Workshop on Case-Based Reasoning,* Washington, D.C.

Goldstone, R. L. & Medin, D. L. (1994a). Similarity, interactive-activation and mapping: In K. J. Holyoak and J. A. Barnden (Eds.), *Advances in Connectionist and Neural Computation Theory, Vol. 2: Connectionist Approaches to Analogy, Metaphor, and Case-Based Reasoning.* Norwood, NJ: Ablex.

Goldstone, R. L., & Medin, D. L. (1994b). Time course of comparison. *Journal of Experimental Psychology: Learning, Memory and Cognition, 20*(1), 29-50.

Goldstone, R. L., Medin, D., & Gentner D. (1991). Relational similarity and the non- independence of features in similarity judgments. *Cognitive Psychology, 23,* 222-264.

Gronlund, S. D., & Ratcliff, R. (1989). The time-course of item and associative information: implications for global memory models. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 15,* 846-858.

Halford (1992). Analogical reasoning and conceptual complexity in cognitive development. *Human development.*

Hammond, K. J. (1986). CHEF: A model of case-based planning. *Proceedings of the Fifth National Conference on Artificial Intelligence* (pp. 267-271). Philadelphia, PA; Los Altos, CA: Morgan Kaufmann.

Hammond, K. (1989). On Functionally motivated vocabularies: An Apologia. In Proceedings: Case-Based Reasoning Workshop, (pp 52-56). San Mateo, CA: Morgan Kaufmann.

Hinton, G. E., & Anderson, J. A. (1989). *Parallel models of associative memory*. Hillsdale, NJ: Erlbaum.
Hintzman, D. L. (1984). MINERVA 2: A simulation model of human memory. *Behavior Research Methods, Instruments, & Computers, 16*, 96-101.

Hintzman, D. (1986). 'Schema abstraction' in a multiple-trace memory model. *Psychological Review, 93*, 411-428.

Hintzman, D. L. (1988). Judgments of frequency and recognition memory in a multiple- trace memory model. *Psychological Review, 95*, 528-551.

Holyoak, K. J., & Koh, K. (1987). Surface and structural similarity in analogical transfer. *Memory & Cognition, 15*, 32-340.

Holyoak, K. J., Novick, L. R., & Melz, E. (1994). Component processes in analogical transfer: Mapping, pattern completion, and adaptation. In K. J. Holyoak & J. A. Barnden (Eds.), *Advances in connectionist and neural computation theory, Vol. 2. connectionist approaches to analogy, metaphor, and case-based reasoning*. Norwood, NJ: Ablex.

Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science, 13*, 295-355.

Hummel, J. E., & Holyoak, K. J. (1992). Indirect analogical mapping. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society* (pp. 516-521). Bloomington, IN.

Kass, A. (1986). Modifying explanations to understand stories. *Proceedings of the Eighth Annual Conference of the Cognitive Science Society* (pp. 691-696). Amherst, MA.

Kass, A. (1989). Strategies for adapting explanations. *Proceedings: Case-Based Reasoning Workshop* (pp. 119-123). San Mateo, CA: Morgan Kaufmann.

Keane, M. (1987). On retrieving analogues when solving problems. *Quarterly Journal of Experimental Psychology, 39A*, 29-41.

Keane, M. (1988a). Analogical mechanisms. *Artificial Intelligence Review, 2*, 229-250.

Keane, M. T. (1988b). *Analogical problem solving*. Chichester: Ellis Horwood (New York: Wiley).

Keane, M. T. (1988c). *Incremental Analogising: Theory & Model* (Technical Rep. No. 38). Human Cognition Research Laboratory, The Open University. Milton Keynes, England.

Keane, M. T. (1991, August). Similarity and ordering constraints. *Proceedings of the Meeting of the Cognitive Science Society.*

Keane, M., & Brayshaw, M. (1988). The incremental analogy machine: a computational model of analogy. *European Session on Machine Learning*. London: Pitman.

Keil, F. C. (1989). *Concepts, kinds, and cognitive development*. Cambridge, MA: MIT Press.

Kettler, B. P., Hendler, J. A., and Anderson, W. A. (1992). Massively parallel support for a case-based planning system. *Proceedings of the NASA Space Applications and Operations Research Workshop* Houston, Texas.

King, J., & Bareiss, R. (1989). Similarity Assessment in Case-Based Reasoning. *Proceedings: Case-Based Reasoning Workshop* (pp. 67-71). San Mateo: Morgan Kaufmann.

Kolodner, J. L. (1984). *Retrieval and organizational structures in conceptual memory: A computer model.* Hillsdale NJ: Erlbaum.

Kolodner, J. L. (1988). *Proceedings of the First Case-Based Reasoning Workshop.* Los Altos, CA: Morgan Kaufmann.

Kolodner, J. L. (1989). Judging which is the "best" case for a case-based reasoner. In *Proceedings: Case-Based Reasoning Workshop,* (pp. 77-81). San Mateo, CA: Morgan Kaufmann.

Kolodner, J. L. (1993) *Case-Based Reasoning*. San Mateo, CA: Morgan Kaufmann.

Kotovsky, L., & Gentner, D. (1990). Pack light: You will go farther. *Proceedings of the Second Midwest Artificial Intelligence and Cognitive Science Society Conference.* Carbondale, IL.

Lange, T. E., & Wharton, C. M. (1992). REMIND: Integrating language understanding and episodic memory retrieval in a connectionist network. *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society.*

Lange, T. E., & Wharton, C. M. (1993). Dynamic memories: Analysis of an integrated comprehension and episodic memory retrieval model. *Proceedings of IJCAI-93.*

Markman, A. G., & Gentner, D. (1990). Analogical mapping during similarity judgments. *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society* (pp. 38-44). Cambridge, MA. Hillsdale, NJ: Erlbaum.

Markman, A. G., & Gentner, D. (1993a). Splitting the differences: A structural alignment view of similarity. *Journal of Memory and Language.*

Markman, A.B., & Gentner, D. (1993b). Structural alignment during similarity comparisons. *Cognitive Psychology, 25,* 431-467.

Martin, C. (1989, May). Indexing using complex features. In *Proceedings of the Case-based Reasoning Workshop* (pp. 26-30).

Medin, D.L., Goldstone, R.L., & Gentner, D. (1993). Respects for similarity. *Psychological Review, 100*(2), 254-278.

Medin, D., & Ortony, A. (1989). Psychological essentialism. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 179-195).

Medin, D. L., & Ross, B. H. (1989). The specific character of abstract thought: Categorization, problem-solving, and induction. In R. J. Sternberg (Ed.), *Advances in the psychology of human intelligence* (Vol. 5, pp. 189-223). Hillsdale, NJ: Erlbaum.

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review, 85,* 207-238.

Minsky, M. (1981). A framework for representing knowledge. In J. Haugland (Ed.), *Mind Design* (pp. 95-128). Cambridge, MA: MIT Press.

Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review, 92,* 289-316.

Novick, L. R. (1988). Analogical transfer: Processes and individual differences. In D. H. Helman (Ed.), *Analogical reasoning: Perspectives of artificial intelligence, cognitive science, and philosophy* (pp. 125-145). Dordrecht, The Netherlands: Kluwer.

Novick, L. R. (1988). Analogical transfer, problem similarity, and expertise. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14,* 510-520.

Pazzani, M. (1989). Indexing strategies for goal-specific retrieval of cases. *Proceedings: Case-based Reasoning Workshop* (pp. 31-35). San Mateo, CA: Morgan Kaufmann.

Porter, B. W. (1989, May). Similarity assessment: Computation vs. representation. In *Proceedings: Case-based Reasoning Workshop* (pp. 82-84). San Mateo, CA: Morgan Kaufmann.

Ratcliff, R., & McKoon, G. (1989). Similarity information versus relational information: Differences in the time course of retrieval. *Cognitive Psychology, 21,* 139-155.

Rattermann, M. J., & Gentner, D. (1987). Analogy and similarity: Determinants of accessibility and inferential soundness. In *Proceedings of the Ninth Annual Meeting of the Cognitive Science Society* (pp. 23-34). Seattle, WA.

Read, S. J., & Cesa, I. L. (1991). This reminds me of the time when...: Expectation failures in reminding and explanation. *Journal of Experimental Social Psychology, 27,* 1-25.

Reed, S. K. (1987). A structure-mapping model for word problems. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13,* 124-139.

Reed, S. K., Ackinclose, C. C., & Voss, A. A. (1990). Selecting analogous problems: similarity versus inclusiveness. *Memory & Cognition, 18* (1), 83-98.

Reed, S. K., Ernst, G. W., & Banerji, R. (1974). The role of analogy in transfer between similar problem states. *Cognitive Psychology, 6,* 436-450.

Reed, L. M., & Weisberg, R. W., (1994) The role of content and abstract information in analogical transfer. *Psychological Bulletin, 115,* (3), 381-400.

Riesbeck, C. K., & Schank, R. C. (1989). *Inside Case-based reasoning.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Rips, L. J. (1989). Similarity, typicality, and categorization. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 21-59). New York: Cambridge University Press.

Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology, 8,* 382-439.

Ross, B. H. (1984). Remindings and their effects in learning a cognitive skill. *Cognitive Psychology, 16,* 371-416.

Ross, B. H. (1987). This is like that: The use of earlier problems and the separation of similarity effects. *Journal of Experimental Psychology: Learning, Memory and Cognition, 13,* 629-639.

Ross, B. H. (1989). Remindings in learning and instruction. In S. Vosniadou & A. Ortony (Eds.), *Similarity and anlogical reasoning* (pp. 438-469). New York: Cambridge University Press.

Schank, R. (1982). *Dynamic Memory.* New York: Cambridge University Press.

Schumacher, R., & Gentner, D. (1987, May). *Similarity-based remindings: The effects of similarity and interitem distance.* Paper presented at the Midwestern Psychological Association, Chicago, IL.

Schumacher, R. M., & Gentner, D. (in preparation). Analogical access: Effects of repetition, competition, and instruction.

Seifert, C. M., McKoon, G., Abelson, R. P., & Ratcliff, R. (1986). Memory connection between thematically similar episodes. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12,* 220-231.

Skorstad, J., Falkenhainer, B., & Gentner, D. (1987). Analogical processing: A simulation and empirical corroboration. *Proceedings of the Sixth National Conference on Artificial Intelligence* (pp. 322-326). Seattle, WA. Los Altos, CA: Morgan Kaufmann.

Skorstad, J., Gentner, D., & Medin, D. (1988). Abstraction processes during concept learning: A structural view. In *Proceedings of the Tenth Annual Conference of the Cognitive Science Society* (pp.419-425). Montreal, Canada.

Smith, L. B. (1984). Young children's understanding of attributes and dimensions: A comparison of conceptual and linguistic measures. *Child Development, 55,* 363-380.

Smolensky, P. (1988). On the proper treatment of connectionism. *Behavior and Brain Sciences, 11,* 1-74.

Spellman, B. A., & Holyoak, K. J. (1992). If Saddam is Hitler then who is George Bush? Analogical mapping between systems of social roles. *Journal of Personality and Social Psychology, 62* (9), 913-933.

Stanfill, C. & Waltz, D. (1986). Toward memory-based reasoning. *Transactions of the ACM, 29* (12), 1213-1228.

Thagard, P., & Holyoak, K. J. (1989). Why indexing is the wrong way to think about analog retrieval. *Proceedings: Case-based Reasoning Workshop* (pp. 36-40). San Mateo, CA: Morgan Kaufmann.

Thagard, P., Holyoak, K. J., Nelson, G., & Gochfeld, D. 1990). Analog retrieval by constraint satisfaction. *Atificial Intelligence, 46,* 259-310.

Van Lehn, K. (1989). *Mind Bugs: the Origins of procedural Misconceptions.* Cambridge, MA: MIT Press.

Waltz, D. (1989, May). Panel discussion on"Indexing Algorithms". *Proceedings: Case-based Reasoning Workshop* (pp. 25-44). San Mateo, CA: Morgan Kaufmann.

Wharton, C. M., Holyoak, K. J., Downing, P. E., Lange, T. E., & Wickens T. D. (1991). Retrieval competition in memory for analogies. In *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society* (pp. 528-533). Hillsdale, NJ: Erlbaum.

Wharton, C. M., Holyoak, K. J., Downing, P. E., Lange, T. E., & Wickens, T. D. (1992). The story with reminding: memory retrieval is influenced by analogical similarity. *Proceeding of the Fourteenth Annual Conference of the Cognitive Science Society.*

Wharton, C. M., Holyoak, K. J., Downing, P. E., Lange, T. E., & Wickens T. D., Melz, E. R. (1994). Below the surface: Analogical similarity and retrieval competition in reminding. *Cognitive Psychology, 26,* 64-101.

Winston, P. H. (1982). Learning new principles from precedents and exercises. *Artificial Intelligence, 19,* 321-350.

GENTNER.TCL 19 MAR 98
ALL_AREA, COMP_MOD, PERCEPTN

Distribution List

Dr. Terry Anderson
Educational Psychology
260C Education Bldg.
University of Illinois
Champaign, IL 61801

Dr. Beth Adelson
Department of Psychology
Rutgers University
Camden, NJ 08102

Director
Human Engineering Lab
ATTN: SLCHE-D
Aberdeen Proving Grounds
MD 21005-5001

Dr. Robert Ahlers
12350 Research Parkway
Human Factors Division, Code 262
Naval Training Systems Center
Orlando, FL 32826

Technical Document Center
AL/HOR-TDC
Wright-Patterson AFB
OH 45433-6503

Dr. David L. Alderton, Code 131
Navy Personnel R&D Center
San Diego, CA 92152-6800

Dr. Terry Allard
Code 1142CS
Office of Naval Research
800 N. Quincy St.
Arlington, VA 22217-5660

Dr. Nancy Allen
Educational Testing Service
Princeton, NJ 08541

Dr. James A. Anderson
Department of Cognitive and
Linguistic Sciences
Brown University
Box 1978
Providence, RI 02912

Dr. John R. Anderson
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. Nancy S. Anderson
Department of Psychology
University of Maryland
College Park, MD 20742

Dr. Phipps Arabie
Graduate School of Management
Rutgers University
92 New Street
Newark, NJ 07102-1895

Technical Director, ARI
5001 Eisenhower Avenue
Alexandria, VA 22333

Edward Atkins
13705 Lakewood Ct.
Rockville, MD 20850

Dr. Courtenay W. Atwell
National Eye Institute
SAVP, Bldg. 31, Rm. 6A47
9000 Rockville Pike
Rockville, MD 20892

Dr. Michael B. Attewell
NYNEX
AI Laboratory
500 Westchester Avenue
White Plains, NY 10604

Dr. Pierre Baldi
Division of Biology, 216-76
California Institute of Technology
Pasadena, CA 91125

prof. dott. Bruno G. Bara
Unità di ricerca di
intelligenza artificiale
Università di Milano
20122 Milano - via F. Sforza 23
ITALY

Dr. William G. Barry
Director of Life and
Environmental Sciences
AFOSR/NL, Nl, Bldg. 410
Bolling AFB, DC 20332-6448

Dr. Thomas G. Bever
Department of Psychology
University of Rochester
River Station
Rochester, NY 14627

Dr. John Black
Teachers College, Box 8
Columbia University
525 West 120th Street
New York, NY 10027

Dr. Michael Blackburn
Code 943
Naval Ocean Systems Center
San Diego, CA 92152-5000

Dr. Colin Blakemore
University Laboratory of Physiology
University of Oxford
Parks Road
Oxford OX1 3PT
UNITED KINGDOM

Dr. Sheila Blumstein
Dean of the College
Brown University
University Hall, Room 218
Providence, RI 02912

Dr. Kenneth R. Boff
AL/CFH
Wright-Patterson AFB
OH 45433-6573

Dr. C. Alan Boneau
Department of Psychology
George Mason University
4400 University Drive
Fairfax, VA 22030

Dr. J. C. Boudreaux
Manufacturing Engineering Lab
National Institute of
Standards and Technology
Gaithersburg, MD 20899

Dr. Gordon H. Bower
Department of Psychology
Stanford University
Stanford, CA 94306

Dr. Robert Breaux
Code 252
Naval Training Systems Center
Orlando, FL 32826-3224

Dr. Jacques Bremond
11 Chemin de la Croix de Buvo
91440 Bures-sur-Yvette
FRANCE

Dr. Gail Carpenter
Center for Adaptive Systems
111 Cummington St., Room 244
Boston University
Boston, MA 02215

Dr. Pat Carpenter
Carnegie-Mellon University
Department of Psychology
Pittsburgh, PA 15213

Dr. Paul R. Chatelier
Perceptronics
1911 North Ft. Myer Dr.
Suite 800
Arlington, VA 22209

Dr. Susan Chipman
Cognitive Science Program
Office of Naval Research
800 North Quincy St.
Arlington, VA 22217-5000

Dr. Reynold B. Christal
UES LAMP Science Advisor
AL/HRMIL
Brooks AFB, TX 78235

Dr. Deborah Clarke
National Institute for Aging
Bldg. 31, Room 5C-35
9000 Rockville Pike
Bethesda, MD 20892

Dr. William J. Clancey
Institute for Research
on Learning
2550 Hanover Street
Palo Alto, CA 94304

Dr. David B. Clement
Department of Psychology
University of South Carolina
Columbia, SC 29208

Chief of Naval Education and
Training (N-5)
NAS Pensacola, FL 32508

Dr. Rodney Cooking
NIMH, Basic Behavior and
Cognitive Science Research
5600 Fishers Lane, Rm 11C-10
Parklawn Building
Rockville, MD 20857

Office of Naval Research
Code 1142
800 N. Quincy Street
Arlington, VA 22217-5000

Library, Code 2
Navy Personnel R&D Center
San Diego, CA 92152-5800

Commanding Officer
Naval Research Laboratory
Code 4827
Washington, DC 20375-5000

Dr. Lynn A. Cooper
Department of Psychology
Columbia University
New York, NY 10027

Dr. Albert T. Corbett
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213

Dr. Kenneth B. Cross
Anacapa Sciences, Inc.
P.O. Box 519
Santa Barbara, CA 93102

CTB/McGraw-Hill Library
2500 Garden Road
Monterey, CA 93940-5380

Dr. Charles E. Davis
Educational Testing Service
Mail Stop 22-T
Princeton, NJ 08541

Dr. Geary Delamete
Bioinformatics
3608 Lyon Street
San Francisco, CA
94123

Dr. Stephanie Doan
University of Illinois

BEST COPY AVAILABLE

Department of Psychology
603 East Daniel Street
Champaign, IL 61820

Defense Technical
Information Center
DTIC/DDA-2
Cameron Station, Bldg 5
Alexandria, VA 22314
(4 Copies)

Dr. Nancy Eldredge
College of Education
Division of Special Education
The University of Arizona
Tucson, AZ 85721

Dr. John Ellis
Navy Personnel R&D Center
Code 13
San Diego, CA 92152-6800

Dr. Jeffrey Elman
University of California,
San Diego
Department of Linguistics, C-008
La Jolla, CA 92093

Dr. George Engelhard, Jr.
Division of Educational Studies
Emory University
210 Fishburne Bldg.
Atlanta, GA 30322

Dr. Carl E. England
Naval Ocean Systems Center
Code 442
San Diego, CA 92152-5000

Dr. Susan Epstein
144 S. Mountain Avenue
Montclair, NJ 07042

ERIC Facility-Acquisitions
1301 Piccard Drive, Suite 300
Rockville, MD 20850-4305

Dr. K. Anders Ericsson
University of Colorado
Department of Psychology
Campus Box 345
Boulder, CO 80309-0345

Dr. Martha Evens
Dept. of Computer Science
Illinois Institute of Technology
10 West 31st Street
Chicago, IL 60616

Dr. Franco Fabbri
Direttore Generale LEONARDO
Piazzale K. Adenauer, 3
00144 ROMA EUR
ITALY

Dr. Beatrice J. Farr
Army Research Institute
PERI-IC
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Marshall J. Farr
Farr-Sight Co.
2520 North Vernon Street
Arlington, VA 22207

Dr. Victor Fields
Department of Psychology
Montgomery College
Rockville, MD 20850

Dr. Leif H. Finkel
Department of Bioengineering
University of Pennsylvania
220 South 33rd Street
Philadelphia, PA 19104-6392

Dr. J. D. Fletcher
Institute for Defense Analyses
888 N. Beauregard St.
Alexandria, VA 22311

Dr. Kenneth D. Forbus
Institute for the Learning
Sciences
Northwestern University
1890 Maple Avenue
Evanston, IL 60201

Dr. Peter T. Fox
Research Imaging Center
Health Science Center at San Antonio
University of Texas
7703 Floyd Curl Drive
San Antonio, TX 78284-7801

Dr. Carl H. Frederiksen
Dept. of Educational Psychology
McGill University
3700 McTavish Street
Montreal, Quebec
CANADA H3A 1Y2

Dr. Alfred R. Fregly
AFOSR/NL, Bldg. 410
Bolling AFB, DC 20332-6448

Dr. Michael Friendly
Psychology Department
York University
Toronto ONT
CANADA M3J 1P3

Dr. Morrill F. Garrett
Director of Cognitive Science
Department of Psychology, Room 312
University of Arizona
Tucson, AZ 85721

Dr. Josh J. Gelfand
Department of Psychology
Princeton University
Princeton, NJ 08544-1010

Dr. Dedre Gentner
Northwestern University
Department of Psychology
2029 Sheridan Road
Swift Hall, Rm 102
Evanston, IL 60208-2710

Chair, Department of
Computer Science
George Mason University
Fairfax, VA 22030

Dr. Helen Gigley
Naval Research Lab., Code 5530
4555 Overlook Avenue, S. W.
Washington, DC 20375-5000

Mr. Matt Gloss
Defense Logistic Agency
Systems Automation Ctr.
DSAC-TMP, Building 27-1
P.O. Box 1605
Columbus, OH 43216-5002

Dr. Dennis Glueckman
National Institute
of Mental Health
Parklawn Bldg.
5600 Fishers Lane
Rockville, MD 20857

Dr. Robert Glaser
Learning Research
& Development Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15260

Dr. Arthur M. Glenberg
University of Wisconsin
W. J. Brogden Psychology Bldg.
1202 W. Johnson Street
Madison, WI 53706

Dr. Sam Glucksberg
Department of Psychology
Princeton University
Princeton, NJ 08544-1010

Prof. Joseph Goguen
PRG, Univ. of Oxford
11 Keble Road
Oxford OX13QD
UNITED KINGDOM

Dr. Susan R. Goldman
Peabody College, Box 45
Vanderbilt University
Nashville, TN 37203

Dr. Timothy Goldsmith
Department of Psychology
University of New Mexico
Albuquerque, NM 87131

Dr. Sherrie Gott
AFHRL/MOMJ
Brooks AFB, TX 78235-5601

Dr. Marilyn E. Greenag
Office of Personnel R&D
1900 E St., NW, Room 6462
Office of Personnel Management
Washington, DC 20415

Dr. Arthur C. Graesser
Department of Psychology
Memphis State University
Memphis, TN 38152

Dr. Wayne Gray
Graduate School of Education
Fordham University
113 West 60th Street
New York, NY 10023

Dr. Bert Green
Johns Hopkins University
Department of Psychology
Charles & 34th Street
Baltimore, MD 21218

Dr. James G. Greeno
School of Education
Stanford University
Room 311
Stanford, CA 94305

Dr. Charles G. Gross
Department of Psychology
Princeton University
Green Hall
Princeton, NJ 08544-1010

Dr. Stephen Grossberg
Center for Adaptive Systems
Room 244
111 Cummington Street
Boston University
Boston, MA 02215

Dr. Gerhard Grossing
Austrian Institute for
Medicinal Studies
Parkgasse 9
Vienna
AUSTRIA A-1030

Naval Command, Control and Ocean
Surveillance Center
RDT&E Division
Attn: Mr. J. Grossmann, Code 44
Bldg. 334
San Diego, CA 92152-5000

BEST COPY AVAILABLE

Dr. Kristian J. Hammond
Department of Computer Science
University of Chicago
1100 E. 58th St.
Chicago, IL 60637

Dr. Ivar Hansen
Forskerin Psykologkjonsette
Akershus/Oslo mil
0015 Oslo 1
NORWAY

Dr. Stephen J. Hanson
Learning & Knowledge
Acquisition Research
Siemens Research Center
755 College Road East
Princeton, NJ 08540

Steven Harnad
Editor, The Behavioral and
Brain Sciences
20 Nassau Street, Suite 240
Princeton, NJ 08542

Dr. Dalwyn Harnish
University of Illinois
51 Gerty Drive
Champaign, IL 61820

Dr. Patrick R. Harrison
Computer Science Department
U.S. Naval Academy
Annapolis, MD 21402-5002

Dr. Barbara Hayes-Roth
Knowledge Systems Laboratory
Stanford University
701 Welch Road, Bldg. C
Palo Alto, CA 94304

Dr. Per Helmersen
University of Oslo
UNIT
Box 1059
0316 Oslo, NORWAY

Ms. Rebecca Hetter
Navy Personnel R&D Center
Code 13
San Diego, CA 92152-6800

Dr. James B. Hoffman
Department of Psychology
University of Delaware
Newark, DE 19711

Dr. Keith Holyoak
Department of Psychology
University of California
Los Angeles, CA 90024

Dr. H. Gene Hondt Van
AFDJ/PW
Administratie Kr. D 364
Van Der Burchlaan 31
Post Box 20702.2500 ES The Hague
The NETHERLANDS

Ms. Julie S. Hough
Cambridge University Press
40 West 20th Street
New York, NY 10011

Dr. William Howell
Chief Scientist
AFHRL/CA
Brooks AFB, TX 78235-5601

Dr. Ben Hudlicka
BBN Laboratories
10 Moulton Street
Cambridge, MA 02254

Dr. Michael P. Huerta
National Institute of
Mental Health, DBSRR
5600 Fishers Lane
Parklawn Building
Rockville, MD 20857

Dr. Earl Hunt
Dept. of Psychology, NI-25
University of Washington
Seattle, WA 98195

Dr. Giorgio Ingargiola
Computer Science Department
Temple University
Philadelphia, PA 19122

Dr. Martin J. Ippel
Center for the Study of
Education and Instruction
Leiden University
P. O. Box 9555
2300 RB Leiden
THE NETHERLANDS

Z. Jacobson
Bureau of Management Consulting
781-340 Laurier Ave., W.
Ottawa, Ontario K1A 5W5
CANADA

Dr. Robert Jannarone
Elec. and Computer Eng. Dept.
University of South Carolina
Columbia, SC 29208

Dr. Kumar Joag-dev
University of Illinois
Department of Statistics
101 Illini Hall
725 South Wright Street
Champaign, IL 61820

Dr. Bonnie E. John
Department of Computer Science
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213

Chair, Department of
Psychology
The Johns Hopkins University
Baltimore, MD 21218

Dr. John Jonides
Department of Psychology
University of Michigan
Ann Arbor, MI 48104

Dr. Marcel Just
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

Dr. J. L. Kaiwi
Code 443DTE
Naval Ocean Systems Center
San Diego, CA 92152-5000

Dr. Michael Kaplan
Office of Basic Research
U.S. Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333-5600

Dr. A. Karmiloff-Smith
MRC-CDU
17 Gordon Street
London
ENGLAND WC1H 0AH

Dr. Milton S. Katz
PSC 802 Box 15
FPO AE 09498-1500

Dr. Steven W. Keele
Department of Psychology
University of Oregon
Eugene, OR 97403

Dr. Douglas Kelly
University of North Carolina
Department of
Statistics, CB #3260
Chapel Hill, NC 27599

Dr. J.A.S. Kelso
Center for Complex Systems
Building MT 9
Florida Atlantic University
Boca Raton, FL 33431

CDR Robert S. Kennedy
Essex Corporation
1040 Woodcock Road
Suite 227
Orlando, FL 32803

Dr. Henry Khachaturian
National Institute of
Mental Health, DBSRR
5600 Fishers Lane
Parklawn Building
Rockville, MD 20857

Dr. David Kieras
Technical Communication Program
TIDAL Bldg., 2360 Bonisteel Blvd.
University of Michigan
Ann Arbor, MI 48109-2108

Dr. Sung-Ho Kim
Educational Testing Service
Princeton, NJ 08541

Dr. Kenneth A. Klivington
The Salk Institute
P.O. Box 85800
San Diego, CA 92186-5800

LCDR David A. Kobus
Code 23
Naval Health Research Center
P. O. Box 85122
San Diego, CA 92186-5122

Dr. Sylvan Kornblum
University of Michigan
Mental Health Research Institute
205 Washtenaw Place
Ann Arbor, MI 48109

Dr. Stephen Kosslyn
Harvard University
1236 William James Hall
33 Kirkland St.
Cambridge, MA 02138

Richard Lamerman
Commandant (G-PWP)
US Coast Guard
2100 Second St., SW
Washington, DC 20593-0001

Dr. Jill Larkin
Carnegie-Mellon University
Department of Psychology
Pittsburgh, PA 15213

Dr. Jill F. Lehman
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213-3890

Dr. Paul E. Lehner
Department of Information
Systems & Engineering
George Mason University
4400 University Drive
Fairfax, VA 22030-4444

Dr. Michael Levine
Educational Psychology
210 Education Bldg.
1310 South Sixth Street
University of IL at
Urbana-Champaign
Champaign, IL 61820-6990

BEST COPY AVAILABLE

Dr. Marcia C. Linn
Graduate School
of Education, EMST
Tolman Hall
University of California
Berkeley, CA 94720

Dr. Margaret S. Livingstone
Department of Neurobiology
Harvard Medical School
220 Longwood Avenue
Boston, MA 02115

Dr. Robert Lloyd
Dept. of Geography
University of South Carolina
Columbia, SC 29208

Logicon Inc. (Attn: Library)
Tactical and Training Systems
Division
P.O. Box 85158
San Diego, CA 92138-5158

Dr. Richard Lowler
ACT
P. O. Box 168
Iowa City, IA 52243

LT(P) C. D. F. Lynn
Command Personnel Applied
Research Coordinator
Maritime Command Headquarters
FMO Halifax, N.S., B3K 2X0
CANADA

Dr. Don Lyon
P. O. Box 2020
Higley, AZ 85236

Dr. Donald MacGregor
Decision Research
1201 Oak St.
Eugene, OR 97401

Vern M. Malec
NPRDC, Code 142
San Diego, CA 92152-6800

Dr. Jane Malin
Mail Code ER22
NASA Johnson Space Center
Houston, TX 77058

Dr. Bruce Mandler
George Mason University
4400 University Drive
Fairfax, VA 22030

Dr. Herb Marg
School of Optometry
University of California
Berkeley, CA 94720

Dr. Elizabeth Martin
AL/HRA, Stop 44
Williams AFB
AZ 85240

Dr. Nadine Martin
Department of Neurology
Center for Cognitive Neuroscience
Temple University School of Medicine
3401 North Broad Street
Philadelphia, PA 19140

Mr. Sten Mattisl
Foreverts Center for Lederskab
Christianshavns Voldgade 8
1424 Kobenhavn K
DENMARK

Dr. Marissa M. Matthews
Department of Computer Science
University of South Carolina
Columbia, SC 29208

Dr. Joseph McLachlan
Navy Personnel Research
and Development Center
Code 14
San Diego, CA 92152-6800

Dr. Michael McNeese
DBT-1, ALICPMI
HLDO 24B
Wright-Patterson AFB, OH 45432

Dr. Alan Meyrowitz
Naval Research Laboratory
Code 5510
4555 Overlook Ave., SW
Washington, DC 20375-5000

Dr. Ryszard S. Michalski
Center for Artificial Intelligence
George Mason University
Science and Tech II, Rm.411
4400 University Drive
Fairfax, VA 22030-4444

Dr. Mortimer Mishkin
Laboratory of Neuropsychology
National Institute of Mental Health
9000 Rockville Pike, Bldg.9, #1N107
Bethesda, MD 20892

Dr. Christine M. Mitchell
School of Indus. and Sys. Eng.
Center for Man-Machine
Systems Research
Georgia Institute of Technology
Atlanta, GA 30332-0205

Dr. J. Anthony Movshon
Center for Neural Science
New York University
6 Washington Place, 8th FL
New York University
New York, NY 10003

Dr. Randy Mumaw
Human Sciences
Westinghouse Science
& Technology Ctr.
1310 Beulah Road
Pittsburgh, PA 15235

Dr. Richard Nakamura
National Institute of
Mental Health, DBBB/CBVRB
5600 Fishers Lane
Parklawn Building
Rockville, MD 20857

Academic Progs. & Research Branch
Naval Technical Training Command
Code N-62
NAS Memphis (75)
Millington, TN 38054

Prof. David Navon
Department of Psychology
University of Haifa
Haifa 31999
ISRAEL

Mr. J. Nelissen
Twente University
Fac. Bibl. Toegepaste Onderwijskunde
P. O. Box 217
7500 AE Enschede
The NETHERLANDS

Director
Training Systems Department
NPRDC (Code 14)
San Diego, CA 92152-6800

Director
Training Technology Department
NPRDC (Code 15)
San Diego, CA 92152-6800

Library, NPRDC

Code 041
San Diego, CA 92152-6800

Librarian
Naval Center for Applied Research
in Artificial Intelligence
Naval Research Laboratory
Code 5510
Washington, DC 20375-5000

Dr. Paul O'Rorke
Information & Computer Science
University of California, Irvine
Irvine, CA 92717

Dr. Stellan Ohlsson
Learning R & D Center
University of Pittsburgh
Pittsburgh, PA 15260

Dr. Judith Reitman Olson
Graduate School of Business
University of Michigan
Ann Arbor, MI 48109-1234

Mathematics Division
Office of Naval Research
Code 1111
800 North Quincy Street
Arlington, VA 22217-5000

Office of Naval Research,
Code 1142CS
800 N. Quincy Street
Arlington, VA 22217-5000
(6 Copies)

Assistant for Training Technology
and Human Factors
Office of the DCNO(MPT) (Op-11B)
Department of the Navy
Washington, DC 20350-2000

Dr. Judith Orasanu
Mail Stop 239-1
NASA Ames Research Center
Moffett Field, CA 94035

Dr. James Ortesky
Institute for Defense Analyses
1801 N. Beauregard St.
Alexandria, VA 22311

Dr. Glenn Osga
NOSC, Code 441
San Diego, CA 92152-6800

Dr. Roy Pea
Institute for the
Learning Sciences
Northwestern University
1890 Maple Avenue
Evanston, IL 60201

Perceptual Science
Code 1142PS
Office of Naval Research
Arlington, VA 22217-5000

Dr. Ray S. Perez
ARI (PERI-II)
5001 Eisenhower Avenue
Alexandria, VA 22333

CDR Frank C. Petho
Naval Postgraduate
School
Code OR/PE
Monterey, CA 93943

Dept. of Administrative Sciences
Code 54
Naval Postgraduate School
Monterey, CA 93943-5026

Dr. Peter Pirolli
School of Education
University of California
Berkeley, CA 94720

BEST COPY AVAILABLE

Prof. Terrence Peggie
Massachusetts Institute
of Technology RES-201
Center for Biological
Information Processing
Cambridge, MA 02139

Dr. Martha Polson
Department of Psychology
University of Colorado
Boulder, CO 80309-0344

Dr. Peter Polson
University of Colorado
Department of Psychology
Boulder, CO 80309-0344

Dr. Joseph Psotka
ATTN: PERI-IC
Army Research Institute
5001 Eisenhower Ave.
Alexandria, VA 22333-5600

Psyc Info - CD and M
American Psychological Assoc.
1200 Uhle Street
Arlington, VA 22201

Dr. Marcus Raichle
Department of Radiation Sciences
Washington University
School of Medicine
510 S. Kingshighway, Box 8225
St. Louis, MO 63110

Mr. Paul S. Rau
Code U-33
Naval Surface Warfare Center
White Oak Laboratory
Silver Spring, MD 20903

Dr. James A. Reggia
Dept. of Computer Science
A. V. Williams Bldg.
University of Maryland
College Park, MD 20742

Dr. Daniel Reisberg
Reed College
Department of Psychology
Portland, OR 97202

Dr. Brian Reiser
Cognitive Science Lab
221 Nassau Street
Princeton University
Princeton, NJ 08542

Dr. Lauren Resnick
Learning R & D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213

Dr. H. M. Reynolds
Chair, Department of
Psychology
Gallaudet College
Kendall Green
Washington, DC 20002

Dr. Gilbert Rinard
Mail Stop K01-14
Grumman Aerospace
Bethpage, NY 11714

Dr. Edwina L. Rissland
Dept. of Computer and
Information Science
University of Massachusetts
Amherst, MA 01003

Dr. Linda G. Roberts
Science, Education, and
Transportation Program
Office of Technology Assessment
Congress of the United States
Washington, DC 20510

Dr. William B. Rouse
Search Technology, Inc.
4725 Peachtree Corners Circle
Suite 200
Norcross, GA 30092

Mr. Louis Roussos
University of Illinois
Department of Statistics
101 Illini Hall
725 South Wright St.
Champaign, IL 61820

Kolonel Drs. R. H. Rozendaal
Commandant Luchtmacht
Elektronische en Technische
School, Postbus 6000
6800 MD ARNHEM
The NETHERLANDS

Dr. Jeffrey D. Schall
Department of Psychology
Vanderbilt University
004 Wilson Building
Nashville, TN 37240

Dr. Mark Schlager
SRI International
333 Ravenswood Ave.
Room BS-131
Menlo Park, CA 94025

Dr. Alan H. Schoenfeld
University of California
Department of Education
Berkeley, CA 94720

Dr. Colleen M. Seifert
Department of Psychology
University of Michigan
330 Packard Road
Ann Arbor, MI 48104

Dr. Terrence J. Sejnowski
Professor
The Salk Institute
P. O. Box 85800
San Diego, CA 92138-9774

Dr. Valerie L. Shalin
Department of Industrial
Engineering
State University of New York
342 Lawrence D. Bell Hall
Buffalo, NY 14260

Mr. Richard J. Shavelson
Graduate School of Education
University of California
Santa Barbara, CA 93106

Mr. Colin Sheppard
Command and Control Dept.
Defence Research Agency
Maritime Division
Portsmouth Hants PO14AA
UNITED KINGDOM

Dr. Randall Shumaker
Naval Research Laboratory
Code 5530
4555 Overlook Avenue, S.W.
Washington, DC 20375-5000

Dr. Edward Silver
LRDC
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15260

Dr. Herbert A. Simon
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. Jerome E. Singer
Department of Medical Psychology
Uniformed Services Univ. of the
Health Sciences
4301 Jones Bridge Road
Bethesda, MD 20814-4799

Dr. Bruce D. Steinberg
Curry College
Milton, MA 02186

Dr. Ted Steinke
Dept. of Geography
University of South Carolina
Columbia, SC 29208

Dr. John Tangney
AFOSR/NL, Bldg. 410
Bolling AFB, DC 20332-6448

Dr. William Uttal
Department of Psychology
Arizona State University
Tempe, AZ 85287-1104

Dr. Jacques Vesonder
University of Geneva
Department of Psychology
Geneva
SWITZERLAND 1204

Dr. Barbara White
School of Education
Tolman Hall, EMST
University of California
Berkeley, CA 94720

Dr. Joseph L. Young
National Science Foundation
Room 320
1800 G Street, N.W.
Washington, DC 20550

BEST COPY AVAILABLE